

UNITED STATES PATENT AND TRADEMARK OFFICE

BEFORE THE PATENT TRIAL AND APPEAL BOARD

META PLATFORMS, INC.; INSTAGRAM, INC.; WHATSAPP LLC; META
PLATFORMS TECHNOLOGIES, LLC; AND GIPHY, INC.

Petitioners

v.

VL COLLECTIVE IP LLC

Patent Owner

U.S. Patent No. 7,266,682

Case IPR2023-00923

**SUPPLEMENTAL DECLARATION OF SYLVIA D. HALL-ELLIS, Ph.D.
RELATING TO HANDLEY (EX1005)**

I, Sylvia D. Hall-Ellis, declare as follows:

I. INTRODUCTION

1. My name is Sylvia D. Hall-Ellis. I have been retained as an expert by Petitioners Meta Platforms, Inc., Instagram, Inc., WhatsApp LLC, Meta Platforms Technologies LLC, and GIPHY, Inc. (“Petitioners”) in this proceeding before the United States Patent and Trademark Office (“USPTO”).

2. I have written this declaration at the request of the Petitioners to provide my expert opinion regarding the authenticity and public availability of a publication. My declaration sets forth my opinions in detail and provides the basis for my opinions regarding the authenticity and public availability of this publication. If called to testify in the above-captioned proceeding, I will testify regarding the opinions and bases set forth below.

3. I reserve the right to supplement or amend my opinions, and bases for them, in response to any additional evidence, testimony, discovery, argument, and/or other additional information that may be provided to me after the date of this declaration.

4. I am being compensated for my time spent working on this matter at my normal consulting rate of \$350 per hour, plus reimbursement for any additional reasonable expenses. My compensation is not in any way tied to the content of this

declaration, the substance of my opinions, or the outcome of this proceeding. I have no other interests in this proceeding or with any of the parties.

5. All the materials that I considered and relied upon are discussed explicitly in this declaration.

II. QUALIFICATIONS

6. I am currently an Adjunct Professor in the School of Information at San José State University

7. A detailed overview of my qualifications is provided in EX1010 at paragraphs 6–10. My qualifications are also explained in more detail in my *curriculum vitae*, which is included as Appendix 1 in previously submitted EX1010.

III. PRELIMINARIES

A. Scope of Declaration and Legal Standards

8. I am not an attorney and will not offer opinions on the law. I am, however, rendering my expert opinion on the authenticity of the documents referenced herein and on when and how each of these documents was disseminated or otherwise made available to the extent that persons interested and ordinarily skilled in the subject matter or art, exercising ordinary diligence, could have located the documents by January 18, 2000, the critical date of U.S. Patent No. 7,266,682 provided to me by Petitioners' counsel.

9. I am informed by Petitioners' counsel that a printed publication

qualifies as publicly accessible as of the date it was disseminated or otherwise made available such that a person interested in and ordinarily skilled in the relevant subject matter could locate it through the exercise of ordinary diligence.

10. While I understand that the determination of public accessibility under the foregoing standard rests on a case-by-case analysis of the facts particular to an individual publication, I also understand that a printed publication is rendered “publicly accessible” if it is cataloged and indexed by a library such that a person interested in the relevant subject matter could locate it (*i.e.*, I understand that cataloging and indexing by a library is sufficient, though there are other ways that a printed publication may qualify as publicly accessible). One manner of sufficient indexing is indexing according to subject matter category. I understand that the cataloging and indexing by a single library of a single instance of a particular printed publication is sufficient, even if the single library is in a foreign country. I understand that, even if access to a library is restricted, a printed publication that has been cataloged and indexed therein is publicly accessible so long as a presumption is raised that the portion of the public concerned with the relevant subject matter would know of the printed publication. I also understand that the cataloging and indexing of information that would guide a person interested in the relevant subject matter to the printed publication, such as the cataloging and indexing of an abstract for the printed publication, is sufficient to render the

printed publication publicly accessible.

11. I understand that routine business practices, such as general library cataloging and indexing practices, can be used to establish an approximate date on which a printed publication became publicly accessible.

B. Persons of ordinary skill in the art.

12. I am told by Petitioners' counsel that the subject matter of this proceeding relates to data transmission, and specifically, a method and system for transmitting data from a transmitter to a receiver.

13. I have been informed by Petitioners' counsel that a "person of ordinary skill in the art at the time of the invention" ("POSITA") is a hypothetical person who is presumed to be familiar with the relevant field and its literature at the time of the inventions. This hypothetical person is also a person of ordinary creativity, capable of understanding the scientific principles applicable to the pertinent field.

14. I have been informed by Petitioners' counsel that Petitioners contend that persons of ordinary skill in this subject matter or art would have had a Bachelor's degree in computer science, electrical engineering, or a related field, and two years of experience; or by obtaining a Master's degree in electrical engineering or a related field, but having only one to two years of experience; or by having no formal education but experience in network data transmission and

computer networking of at least eight years. Additional education might compensate for less experience, and vice-versa.

15. Regardless of whether Petitioners' or Patent Owner's definition of a person of ordinary skill in the art is used, it is my opinion that such a person would have been engaged in research, learning, study, and practice in the field, and formal instruction so that bibliographic resources relevant to his or her research would be familiar. Before January 18, 2000, such a person would have had access to a vast array of long-established print resources in data transmission and computer networking as well as to a rich set of online resources providing indexing information, abstracts, and full text services.

IV. LIBRARY PROFESSIONAL PRACTICES

16. In preparing this declaration, I used authoritative databases, such as the OCLC bibliographic database, the Library of Congress Online Catalog, and the Library of Congress Subject Authorities, to confirm citation details of the various publications discussed. Unless I note otherwise below in reference to a specific serial publication, it is my expert opinion that this standard protocol was followed for the serial publication discussed below.

17. *Indexing.* A researcher may discover material relevant to his or her topic in a variety of ways. One common means of discovery is to search for relevant information in an index of periodicals and other publications. Having

found relevant material, the researcher will then normally obtain it online, look for it in libraries, or purchase it from the publisher, a bookstore, a document delivery service, or other provider. Sometimes, the date of a document's public accessibility will involve both indexing and library date information. However, date information for indexing entries is often unavailable. This is especially true for online indices.

18. Indexing services use a wide variety of controlled vocabularies to provide subject access and other means of discovering the content of documents. The Library of Congress Subject Authorities includes standard forms of terms and cross references that are included in bibliographic records. Subject headings are terms that an individual seeking a document regardless of format can enter in the search bar of the online catalog. Subjects also connect an authenticated term (one included in the Library of Congress subject headings list) with related, broader, and narrower terms. The formats in which these access terms are presented vary from service to service.

19. Online indexing services commonly provide bibliographic information, abstracts, and full-text copies of the indexed publications, along with a list of the documents cited in the indexed publication. These services also often provide lists of publications that cite a given document. A citation of a document is evidence that the document was publicly available and in use by researchers no later than the publication date of the citing document. Prominent indexing services

include *Scopus*, the *IEEE Xplore* database, the *ACM Digital Library*, *Google Scholar*, and the *Internet Archive*.

V. ONLINE CATALOGS AND MARC RECORDS

20. MARC is a framework into which descriptive bibliographic data are transcribed to interact with the software in online library catalogs to provide access to books, journals, and other resources in the collection. The bibliographic data provide points of access and can be searched by a POSITA to identify and obtain resources in the library collection. An information seeker (or POSITA) can search a local online library catalog or the holdings of a group of libraries in a state or region or in the global catalog WorldCat.

21. MARC records are not designed for public viewing. Although a significant number of libraries provide access to the MARC version of a bibliographic record, the public display is designed to show information in a succinct manner that is quickly understood and useful to the information seeker. Libraries determine the default search for the online catalog to make the entry of search terms efficient and result in a successful search. Information seekers can enter a keyword, title, author, or standard number for the item. Libraries may also provide a search capability called “Summon” that allows the information seeker to enter known information about the item to conduct a search.

22. Learning to use an online library catalog begins in elementary school

and continues through high school and post-secondary studies. Undergraduate students further develop their skills to use the online catalog at the academic institution where they are studying through information literacy instruction which is provided by professional reference and user experience librarians. The skills required to search, identify, and obtain the resources identified in this declaration/report are fundamental to using the online catalogs in libraries and information centers. Therefore, a POSITA who has attained the undergraduate and/or graduate educational levels described above (see paragraph 17) possesses the searching and retrieval skills needed to search library catalogs.

23. Professional catalog and metadata librarians prepare and enhance MARC records to provide points of access for each item in the library collection. Using a set of internationally adopted guidelines, rules for the assignment of classification numbers, and the assignment of subject headings from standardized lists, taxonomies, folksonomies, and locally developed lists. An information seeker (including POSITAs) can search the online library catalog in a number of ways to identify a specific item or resources on a defined topic.

24. The continued development of the software that supports integrated online library catalogs resulted in the potential to search by an individual creator or author, title proper or variant title, a series that contains the desired item, subjects, corporate bodies responsible in part for the creation of an item, or standard

identification numbers associated with the item (e.g., International Standard Book Numbers (ISBN), International Standard Serial Numbers (ISSN), and Publisher's Numbers). Each of these data elements appears in a prescribed location within the MARC framework so that the software can parse, index, and display data elements resulting from a user-initiated search.

25. Examining the MARC records for a specific item reveals the comprehensive data transcribed about a particular item at the time that cataloging and classification occurred. In addition to the creator, title, subjects, and standard numbers, additional information may provide additional and relevant data depending on the type of resource. Understanding the full extent of bibliographic data for an item and the points of access associated with it provide essential information that can be used to determine the indexing and public availability for documents described in this declaration/report.

VI. LIBRARY CATALOGING PRACTICES

26. I am fully familiar with the library cataloging standard known as the MARC standard, which is an industry-wide standard method of storing and organizing library catalog information.¹ MARC was first developed in the 1960's by the Library of Congress. A MARC-compatible library is one that has a catalog

¹ The full text of the standard is available from the Library of Congress at <http://www.loc.gov/marc/bibliographic/> (last visited April 20, 2023).

consisting of individual MARC records for each of its items. Today, MARC is the primary communications protocol for the transfer and storage of bibliographic metadata in libraries.²

27. Since at least the early 1970s and continuing to the present day, MARC has been the primary communications protocol for the transfer and storage of bibliographic metadata in libraries.³ As explained by the Library of Congress:

You could devise your own method of organizing the bibliographic information, but you would be isolating your library, limiting its options, and creating much more work for yourself. Using the MARC standard prevents duplication of work and allows libraries to better share bibliographic resources. Choosing to use MARC enables libraries to acquire cataloging data that is predictable and reliable. If a library were to develop a “home-grown” system that did not use MARC records, it would not be taking advantage of an

² Almost every major library in the world is MARC-compatible. *See, e.g., MARC Frequently Asked Questions (FAQ)*, Library of Congress, <https://www.loc.gov/marc/faq.html> (last visited April 20, 2023) (“MARC is the acronym for Machine-Readable Cataloging. It defines a data format that emerged from a Library of Congress-led initiative that began nearly fifty years ago. It provides the mechanism by which computers exchange, use, and interpret bibliographic information, and its data elements make up the foundation of most library catalogs used today.”). MARC is the ANSI/NISO Z39.2-1994 (reaffirmed 2016) standard for Information Interchange Format.

³ A complete history of the development of MARC can be found in *MARC: Its History and Implications* by Henrietta D. Avram (Washington, DC: Library of Congress, 1975) and available online from the Hathi Trust (<https://babel.hathitrust.org/cgi/pt?id=mdp.39015034388556;view=1up;seq=1>; last visited April 20, 2023).

industry-wide standard whose primary purpose is to foster communication of information.

Using the MARC standard also enables libraries to make use of commercially available library automation systems to manage library operations. Many systems are available for libraries of all sizes and are designed to work with the MARC format. Systems are maintained and improved by the vendor so that libraries can benefit from the latest advances in computer technology. The MARC standard also allows libraries to replace one system with another with the assurance that their data will still be compatible.

Why Is a MARC Record Necessary? LIBRARY OF CONGRESS.⁴

28. Thus, almost every major library in the world is MARC-compatible. See, e.g., *MARC Frequently Asked Questions (FAQ)*, LIBRARY OF CONGRESS.⁵ (“MARC is the acronym for MACHine-Readable Cataloging. It defines a data format that emerged from a Library of Congress-led initiative that began fifty years ago. It provides the mechanism by which computers exchange, use, and interpret bibliographic information, and its data elements make up the foundation of most library catalogs used today.”). MARC is the ANSI/NISO Z39.2-1994 standard (reaffirmed in 2016) for Information Interchange Format. The full text of the

⁴ <http://www.loc.gov/marc/umb/um01to06.html#part2>

⁵ <https://www.loc.gov/marc/faq.html>

standard is available from the Library of Congress.⁶

29. A MARC record comprises several fields, each of which contains specific data about the work. Each field is identified by a standardized, unique, three-digit code corresponding to the type of data that follow. For example, a work's title is recorded in Field 245; the primary author of the work is transcribed in Field 100; an item's International Standard Book Number ("ISBN") consisting of ten or thirteen digits is transcribed in Field 020; an item's International Standard Serial Number ("ISSN") is transcribed in Field 022; the Library of Congress classification notation is recorded in Field 050; and the publication date is recorded in Field 260 under the subfield "c." If a work is a periodical, then its publication frequency is recorded in Field 310, and the publication dates (e.g., the first and last publication) are recorded in Field 362, which is also referred to as the enumeration/chronology field.⁷

30. The library that created the record is recorded in Field 040 in subfield "a" with a unique library code. When viewing the MARC record online via Online Computer Library Center's ("OCLC") bibliographic database, hovering over this code with the mouse reveals the full name of the library. I used this method of "mousing over" the library codes in the OCLC database to identify the originating

⁶ <http://www.loc.gov/marc/bibliographic/>

⁷ <http://www.loc.gov/marc/bibliographic/bd3xx.html>

library for the MARC records discussed in this declaration. Where this “mouse over” option was not available, I consulted the Directory of OCLC Libraries to identify the institution that created the MARC record.⁸

31. MARC records also include several fields that include subject matter classification information. An overview of MARC record fields is available through the Library of Congress.⁹ For example, 6XX fields are termed “Subject Access Fields.”¹⁰ Among these, for example, is the 650 field; this is the “Subject Added Entry – Topical Term” field.¹¹ The 650 field is a “[s]ubject added entry in which the entry element is a topical term.” These entries “are assigned to a bibliographic record to provide access according to generally accepted thesaurus-building rules (e.g., *Library of Congress Subject Headings* (LCSH), *Medical Subject Headings* (MeSH)).” *Id.* Further, MARC records include call numbers, which themselves include a classification number. For example, the 050 field is the “Library of Congress Call Number.”¹² A defined portion of the Library of Congress Call (LCC) Number is the classification number, and “source of the classification number is *Library of Congress Classification* and the *LC*

⁸ <https://www.oclc.org/en/contacts/libraries.html>

⁹ <http://www.loc.gov/marc/bibliographic/>

¹⁰ <http://www.loc.gov/marc/bibliographic/bd6xx.html>

¹¹ <http://www.loc.gov/marc/bibliographic/bd650.html>

¹² <http://www.loc.gov/marc/bibliographic/bd050.html>

Classification-Additions and Changes.” Thus, included in the 050 field is a subject matter classification. Further, the 082 field is the “Dewey Decimal Call Number.”¹³ A defined portion of the Dewey Decimal Call (DDC) Number is the classification number, and “source of the classification number is the *Dewey Decimal Classification and Relative Index.*” Thus, included in the 082 field is a subject matter classification. Each item in a library has a single classification number. A library selects a classification scheme (*e.g.*, the Library of Congress classification scheme just described or a similar scheme such as the Dewey Decimal classification scheme) and uses it consistently. When the Library of Congress assigns the LCC classification number, it appears as part of the 050 field. When the Library of Congress assigns the DDC classification number, it appears as part of the 082 field. If a local library assigns the classification number, it appears in a 090 field. In either scenario, the MARC record includes a classification number that represents a subject matter classification.

32. The OCLC was created “to establish, maintain and operate a computerized library network and to promote the evolution of library use, of libraries themselves, and of librarianship, and to provide processes and products for the benefit of library users and libraries, including such objectives as increasing

¹³ <http://www.loc.gov/marc/bibliographic/bd082.html>

availability of library resources to individual library patrons and reducing the rate of rise of library per-unit costs, all for the fundamental public purpose of furthering ease of access to and use of the ever-expanding body of worldwide scientific, literary and educational knowledge and information.”¹⁴ Among other services, OCLC and its members are responsible for maintaining the WorldCat database (<http://www.worldcat.org/>), used by independent and institutional libraries throughout the world.

33. OCLC also provides its members online access to MARC records through its OCLC bibliographic database. When an OCLC member institution acquires a work, it creates a MARC record for this work in its computer catalog system in the ordinary course of its business. MARC records created at the Library of Congress are tape-loaded into the OCLC database through a subscription to MARC Distribution Services daily or weekly. Once the MARC record is created by a cataloger at an OCLC member institution or is tape-loaded from the Library of Congress, the MARC record is then made available to any other OCLC members online, and therefore made available to the public. Accordingly, once the MARC record is created by a cataloger at an OCLC member institution or is tape-loaded from the Library of Congress or another library anywhere in the world, any

¹⁴ Third Article, Amended Articles of Incorporation of OCLC Online Computer Library Center, Incorporated (available at <https://www.oclc.org/content/dam/oclc/membership/articles-of-incorporation.pdf>).

publication corresponding to the MARC record has been cataloged and indexed according to its subject matter such that a person interested in that subject matter could, with reasonable diligence, locate and access the publication through any library with access to the OCLC bibliographic database or through the Library of Congress.

34. When an OCLC member institution creates a new MARC record, OCLC automatically supplies the date of creation for that record. The date of creation for the MARC record appears in the fixed Field (008), characters 00 through 05. The MARC record creation date reflects the date on which the item was first acquired or cataloged. Initially, Field 005 of the MARC record is automatically populated with the date the MARC record was created in year, month, day format (YYYYMMDD) (some of the newer library catalog systems also include hour, minute, second (HHMMSS)). Thereafter, the library's computer system may automatically update the date in Field 005 every time the library updates the MARC record (*e.g.*, to reflect that an item has been moved to a different shelving location within the library). Field 005 is visible when viewing a MARC record via an appropriate computerized interface, but when a MARC record is printed to hardcopy, no "005" label appears. The initial Field 005 date (*i.e.*, the date the MARC record was created) does appear, however, next to the

label “Entered.”¹⁵ The date upon which the most recent update to Field 005 occurred also appears, next to the label “Replaced.” Thus, when an item’s MARC record has been printed to hardcopy—as is the case with the exhibits to this declaration—the date reflected next to the label “Entered” is necessarily on or after the date the library first cataloged and indexed the underlying item.

35. Once one library has cataloged and indexed a publication by creating a MARC record for that publication, other libraries that receive the publication do not create additional MARC records—the other libraries instead rely on the original MARC record. They may update or revise the MARC record to ensure accuracy, but they do not replace or duplicate it. This practice does more than save libraries from duplicating labor. It also enhances the accuracy of MARC records. Further, it allows librarians around the world to know that a particular MARC record is authoritative (in contrast, a hypothetical system wherein duplicative records were created would result in confusion as to which record is authoritative).

36. Catalogers can create MARC records for all types of print, online, and digital resources. The date of creation of the MARC record by a cataloger at an OCLC member institution reflects when the underlying item is accessible to the

¹⁵ Field 005 is visible when viewing a MARC record via an appropriate computerized interface. But when a MARC record is printed directly to hardcopy from the OCLC database, the “005” label is not shown. The date in the 005 field instead appears next to the label “Replaced.”

public. Upwards of two-thirds to three-quarters of book sales to libraries come from a jobber or wholesaler for online and print resources. These resellers make it their business to provide books to their customers as fast as possible, often providing turnaround times of only a single day after publication. Libraries purchase a significant portion of the balance of their books directly from publishers themselves, which provide delivery on a similarly expedited schedule. In general, libraries make these purchases throughout the year as the books are published and shelve the books as soon thereafter as possible to make the books available to their patrons. Thus, books are generally available at libraries across the country within just a few days of publication.

VII. DOCUMENT 1: ATTACHMENT 1J (“HANDLEY”)

37. Attached hereto as Attachment 1J is a copy of the journal article titled “Very Large Conferences on the Internet: the Internet Multimedia Conferencing Architecture” by Mark Handley, Jon Crowcroft, C. Bormann, and J. Ott (hereafter “Handley”). The Handley article was published in Volume 31, Issue 3 of the journal titled *Computer Networks* and was found in the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison. The Handley article appears beginning on page 191 of this issue dated February 1999. Attachment 1J is a true and correct copy of the Handley article (pages 191-204) and also includes a copy of the *Computer Networks* cover pages and table of contents. Specifically, the text

of the article is complete; no pages are missing, and the text on each page appears to flow seamlessly from one page to the next; further, there are no visible alterations to the document. Attachment 1J was found within the custody of a library – a place where an authentic copy of this journal would likely be. Attachment 1J is a true and correct copy in a condition that creates no suspicion about its authenticity.

38. The Handley article is also available from the publisher, Elsevier,¹⁶ the *ACM Digital Library*,¹⁷ and the digital repositories *ResearchGate*¹⁸ and *Semantic Scholar*.¹⁹ The publisher's website indicates that the Handley article was publicly available online on June 9, 1999 (see Attachment 1A); the *ACM Digital Library*'s website indicates that the Handley article was publicly available on February 11, 1999 (see Attachment 1B).

39. In the early years of the twenty-first century, publishers began offering digital versions of journals to subscribers of the print versions. The rapid

¹⁶ <https://www.sciencedirect.com/science/article/abs/pii/S0169755298002621>

¹⁷ <https://dl.acm.org/doi/10.1016/S0169-7552%2898%2900262-1>

¹⁸

https://www.researchgate.net/publication/228877262_The_internet_multimedia_conferencing_architecture

¹⁹ <https://www.semanticscholar.org/paper/The-Internet-Multimedia-Conferencing-Architecture-Handley-Ott/0c9fea4c1baeca2788720e36354d08be205ce7fe>

expansion of the Internet provided an information sharing environment which publishers used to deliver content to subscribers more rapidly; a venue for the development of private digital repositories (*i.e.*, *ResearchGate* and *Semantic Scholar*) that offered permanent locations to access documents; and an accelerated method for researchers and faculty members to connect with colleagues globally and share their findings and new knowledge. Creators of new intellectual content (*i.e.*, authors, editors, compilers) continued to submit their manuscripts of papers, journal articles, and books to well-known, recognized publishers with the understanding that a digital version could precede the print version by weeks or months. The acceleration of public availability of journal articles, books, and other documents became the new norm as digital versions were released sooner than the print versions. It was further common practice for authors to upload versions as submitted to the publishers to their personal or academic website. These uploads often were versions that were not edited by professional editors or publishers and may include typographical or grammatical or other non-substantive errors.

40. Attached hereto as Attachment 1C is a true and correct copy of the MARC record for the journal titled *Computer Networks* for the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison. Library ownership is indicated by the presence of the library’s code (“GZM”) in field 049. The library continues to update this MARC record and enhanced the MARC record to meet

current cataloging rules. The most recent enhancement to Attachment 1C occurred on April 9, 2015, as shown in field 005 (“20150409”). I personally identified and retrieved the MARC record that is Attachment 1C.

41. Based on finding a print copy of Attachment 1J in the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison and MARC record in its online library catalog attached as Attachment 1C, it is my opinion that the Handley article published in the journal titled *Computer Networks* was publicly available on February 11, 1999, as shown on the *ACM Digital Library* website.

42. Attachment 1C further includes an entry in field 050 (“TK5101.5 \$b .C647”)—as described above, a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification system). Attachment 1C also shows that Attachment 1J was catalogued with four descriptor terms reading “Computer networks \$v Periodicals” (see Attachment 1D, Library of Congress subject heading sh2008101495), “Integrated digital services networks \$v Periodicals” (see Attachment 1E, Library of Congress subject heading sh86000852 and Attachment 1F, Library of Congress subject heading sh85099890), “Telecommunication systems \$v Periodicals” (see Attachment 1G, Library of Congress subject heading sh2010115961), and “Digital communications \$v Periodicals” (see Attachment 1H, Library of Congress subject heading sh85037972 and Attachment 1F, Library of Congress subject heading

sh85099890) in the 650 fields.

43. As noted in the holdings information (362 field), the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison has received the journal titled *Computer Networks* since January 1999 and continues to receive the publication.²⁰ In view of the MARC record for Attachment 1J, the Handley article was publicly available on February 11, 1999, because the serial title had been received, cataloged, and indexed in the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison and made part of its online catalog database.

44. Attached hereto as Attachment 1I is a true and correct copy of the MARC record for the journal titled *Computer Networks* obtained from the OCLC bibliographic database. As previously noted, the library that created the record is recorded in field 040 with a unique library code. For Attachment 1I, that library code is “UHC,” which means that the MARC record for this serial was cataloged at the University of Houston – Clear Lake Library. As can be seen in the “Entered” field in the MARC record for this exhibit, a cataloger at the University of Houston – Clear Lake created OCLC record number 40811573 on February 17, 1999, as shown in the “Entered” field (“19990217”). The library continues to update this

²⁰ The publication was originally titled the *Computer Networks and ISDN Systems* (January 1976-December 1998) and adopted the current title, *Computer Networks*, in January 1999.

MARC record and enhanced the MARC record to meet current cataloging rules. The most recent enhancement to Attachment 1I occurred on March 27, 2019, as shown in the “Replaced” field (“20190327”). I personally identified and retrieved the MARC record that is Attachment 1I. The “BLvl” entry in Attachment 1I is “s,” which indicates that the journal titled *Computer Networks* is a serial publication. Field 310 of Attachment 1I reads “Monthly.” Accordingly, the MARC record for Attachment 1J corresponds to the journal titled *Computer Networks* from the time the serial title adopted the title and includes the date of the Handley article.

45. Attachment 1I includes an entry in field 050 (“TK5105.5 \$b .C647”)—as described above, a subject matter classification number consistent with the Library of Congress classification system (analogous to the Dewey Decimal classification system). Attachment 1I further includes four descriptor terms reading “Computer networks \$v Periodicals” (see Attachment 1D, Library of Congress subject heading sh2008101495), “Integrated digital services networks \$v Periodicals” (see Attachment 1E, Library of Congress subject heading sh86000852 and Attachment 1F, Library of Congress subject heading sh85099890), “Telecommunication systems \$v Periodicals” (see Attachment 1G, Library of Congress subject heading sh2010115961), and “Digital communications \$v Periodicals” (see Attachment 1H, Library of Congress subject heading sh85037972 and Attachment 1F, Library of Congress subject heading sh85099890) in the 650

fields. Thus, as of its cataloging, the publication corresponding to the MARC record attached hereto as Attachment 1I was indexed according to its subject matter by virtue of at least two independently sufficient classifications: the field 050 entry and the field 650 entries. As of February 17, 1999, the MARC record attached hereto as Attachment 1I was accessible through any library with access to the OCLC bibliographic database or the online catalog at a library that subscribed to the serial, which means that the corresponding publication was publicly available on or before that same date through any library with access to the OCLC bibliographic database or through an individual library. Attachment 1I indicates that the journal titled *Computer Networks* as cataloged at the University of Houston – Clear Lake Library is currently available from 185 libraries.

46. While Attachment 1J is a copy of the Handley article (as well as the *Computer Networks* cover pages and table of contents) which was obtained from the Karl F. Wendt Engineering Library at the University of Wisconsin – Madison, Attachment 1K is a copy of the Handley article which was downloaded from the publisher’s website. After comparing Attachment 1J with Attachment 1K, it is my opinion that the Handley articles, included in Attachment 1J and Attachments 1K, contain identical text. Each is a true and accurate copy of the journal article and claims what Petitioner purports it to be. There are no missing pages.

47. In view of the above, this issue of the journal titled *Computer*

Networks was publicly available by at least February 11, 1999, because by that date it had been published online by the *ACM Digital Library*. For these reasons, it is my opinion that Attachment 1J was published and accessible to the public by at least February 11, 1999.

VIII. SUMMARY OF OPINIONS

48. In view of the foregoing, it is my opinion that the publications described above were publicly available no later than the corresponding date listed in the table below:

Exhibit	Publication	Publicly Available No Later Than
Attachment 1J	M. Handley, J. Crowcroft, C. Bormann, J. Ott, <i>Very Large conferences on the Internet: the Internet multimedia conferencing architecture</i> , 31 COMPUTER NETWORKS, 191-204 (February 11, 1999)	February 11, 1999

49. I declare that all statements made herein of my knowledge are true, and that all statements made on information and belief are believed to be true, and that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Dated: January 17, 2024

By: 
Sylvia D. Hall-Ellis, Ph.D.

ATTACHMENT 1A



Very large conferences on the Internet: the Internet multimedia conferencing architecture

M. Handley^a , J. Crowcroft^b, C. Bormann^c, J. Ott^c

^a Information Sciences Institute, University of Southern California, Marina Del Rey, USA

^b Department of Computer Science, University College London, London, UK

^c Informatics Department, University of Bremen, Bremen, Germany

Available online 9 June 1999.

[Show less](#) 

 Share  Cite

[https://doi.org/10.1016/S0169-7552\(98\)00262-1](https://doi.org/10.1016/S0169-7552(98)00262-1) 

[Get rights and content](#) 

Abstract

In this paper we provide an overview of multimedia conferencing on the Internet. The protocols mentioned are all specified elsewhere as Internet-Drafts or RFCs. Each RFC gives details of the protocol itself, how it works and what it does. This document attempts to provide the reader with an overview of how the components fit together and of some of the assumptions made, as well as some statement of direction for those components still in a nascent stage.

Introduction

The Internet is not currently very good at carrying audio and video. This is hardly surprising as it was not designed or engineered with real-time traffic in mind, but there has recently been a great deal of interest in using the Internet for telephony services. Part of this has come from pricing anomalies that make Internet telephony somewhat artificially cheaper than traditional telephone services, but this is not the whole story. The Internet itself is improving to better handle traffic such as audio and video, and in the medium term, the Internet should be able to provide good quality realtime multimedia services, although such quality improvements are likely to incur additional charges.

However, the real interest in using the Internet for audio and video should come from the prospect for a single ubiquitous communications network that not only allows traditional telephony services, but also video, shared collaboration tools, and through IP Multicast, multi-party conferences and multimedia sessions that scale from small group meetings through to television-sized audiences. In principle, this may lead to a “democratization” of telecommunication services, where licenses to broadcast are not required to control physical access to the limited

Meta 1018
Meta v. VL Collective
IPR2023-00923

broadcast medium (although they may still be required for political reasons).

It is far from clear what services will eventually emerge using such communications capabilities. We can only say that the technical capability to have large numbers of sessions ranging in audience from hundreds to millions of participants, largely unlimited by geographic boundaries, will lead to services and social structures that do not exist today. However, we can describe the basic technologies that are likely to bring about such changes, and in this paper we attempt to provide such an overview. We leave it to the reader to imagine the uses to which this technology will be put.

Section snippets

The technology

In conjunction with computers, the term “conferencing” is often used in two different ways: firstly, to refer to the *asynchronous* exchanges of messages between multiple users, as with bulletin boards or mailing lists; secondly, to refer to *synchronous* or so-called “real-time” conferencing, including audio, video, shared whiteboards and other applications. This paper is about the architecture for this latter application in an Internet environment.

There are other infrastructures for conferencing...

Multicast traffic distribution

IP multicast provides efficient many-to-many data distribution in an Internet environment. It is easy to view IP multicast as simply an optimisation for data distribution, and indeed this is the case, but IP multicast can also result in a different way of thinking about application design. To see why this might be the case, examine the IP multicast service model, as described by Van Jacobson [9]:

- Senders just send to *the group*....
- Receivers express an interest in receiving data sent to *the group*....
- ...

...

Internet service models

Traditionally the Internet has provided so-called *best-effort* delivery of datagram traffic from senders to receivers. No guarantees are made regarding when or if a datagram will be delivered to a receiver, however datagrams are normally only dropped when a router exceeds a queue size limit due to congestion. The best-effort Internet service model does not assume FIFO queuing, although many routers have implemented this.

With best-effort service, if a link is not congested, queues will not build...

Transport protocols

So-called real-time delivery of traffic requires little in the way of transport protocol. In particular, real-time traffic that is sent over more than trivial distances is not retransmittable.

With packet multimedia data there is no need for the different media comprising a conference to be carried in the same packets. In fact it simplifies receivers if different media streams are carried in separate flows (i.e., separate

ports and/or separate multicast groups). This also allows the...

Conference control

Conferences come in many shapes and sizes, but there are only really two models for conference control: light-weight sessions and tightly coupled conferencing. For both models, rendezvous mechanisms are needed. Note that the conference control model is orthogonal to issues of quality of service and network resource reservation, and it is also orthogonal to the mechanism for discovering the conference [25].

Light-weight sessions are multicast based multimedia conferences that lack explicit...

Conference setup

There are two basic forms of conference discovery mechanism. These are session advertisement and session invitation. Session advertisements are provided using a session directory, and inviting a user to join a session is provided using a session invitation protocol such as SIP 12, 26 or H.323 [17]....

Security

There is a temptation to believe that multicast is inherently less private than unicast communication since the traffic visits so many more places in the network. In fact, this is not the case except with broadcast and prune type multicast routing protocols [5]. However, IP multicast does make it simple for a host to anonymously join a multicast group and receive traffic destined to that group without the other senders' and receivers' knowledge. If the application requirement (conference...

Summary

This document is an attempt to gather together in one place the set of assumptions behind the design of the Internet Multimedia Conferencing architecture, and the services that are provided to support it.

Fig. 6 shows an example time sequence involved in setting up a light-weight session between two sites. In this case, site A creates a session advertisement, and some time later starts sending a media stream even though there may be no receiver at that time. Some time later, site B joins the...

Acknowledgements

Acknowledgments are due to the End-to-End Research Group, the Int-serv, RSVP, MMUSIC and AVT working groups of the IETF, and discussion with colleagues at UCL. The earliest clear exposition of some of the ideas here was presented at ACM SIGCOMM 1994 in London by Van Jacobson....

References (29)

- A. Ballardie, P. Francis, J. Crowcroft, An architecture for scalable inter-domain multicast Routing, ACM SIGCOMM 1993,...
- C. Bormann, Providing Integrated Services over Low-bitrate Links, Internet-Draft draftietf-issll-isslow-03.txt, Work in...
- CCITT (Consultative Committee on International Telegraphy and Telephony), Recommendation X.509: The Directory -...

W. Dang, Reliable file transfer in the multicast domain, University of Hawaii Technical Report, August...

S. Deering, C. Partridge, D. Waitzman, Distance Vector Multicast Routing Protocol, IETF RFC 1075, November...

S. Deering, Multicast routing in inter-networks and extended LANs, ACM SIGCOMM 88, August 1988, pp. 55–64 and Host...

S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.-G. Liu, L. Wei, An architecture for wide area multicast routing,...

Estrin, Farinacci, Helmy, Thaler, Deering, Handley, Jacobson, Liu, Sharma, Wei, Protocol Independent Multicast-Sparse...

S. Floyd, V. Jacobson, S. McCanne, C.-G. Liu, L. Zhang, A reliable multicast framework for light-weight sessions and...

M. Handley, J. Crowcroft, Network Text Editor (NTE): A scalable shared text editor for the MBone, Proc. ACM SIGCOMM 97,...

 View more references

Cited by (21)

[Multiparty Video Conferencing over Internet ↗](#)

2018, VoIP Handbook: Applications, Technologies, Reliability, and Security

[Reflections on security options for the real-time transport protocol framework ↗](#)

2014, Proceedings - International Conference on Network Protocols, ICNP

[Session and media signalling for communication components-based open multimedia conferencing Web service over IP networks ↗](#)

2011, European Transactions on Telecommunications

[Multiparty audioconferencing on wireless networks ↗](#)

2008, Broadband Mobile Multimedia: Techniques and Applications

[P2PConf: A medium-size P2P internet conference with effective floor control ↗](#)

2008, 2008 International Conference on Information Networking, ICOIN

[Network Convergence: Services, Applications, Transport, and Operations Support ↗](#)

2007, Network Convergence: Services, Applications, Transport, and Operations Support

 View all citing articles on Scopus ↗

Recommended articles (6)

Research article

Meta 1018
Meta v. VL Collective
IPR2023-00923

[Contrasting patterns of fatty acid composition and oil accumulation during fruit growth in several olive varieties and locations in a non-Mediterranean region](#)

European Journal of Agronomy, Volume 52, Part B, 2014, pp. 237-246

[Show abstract](#) ✓

Research article

[Hyperspectral FTIR imaging of olive fruit for understanding ripening processes](#)

Postharvest Biology and Technology, Volume 145, 2018, pp. 74-82

[Show abstract](#) ✓

Research article

[Generic models for rapid detection of vanillin and melamine adulterated in infant formulas from diverse brands based on near-infrared hyperspectral imaging](#)

Infrared Physics & Technology, Volume 116, 2021, Article 103745

[Show abstract](#) ✓

Research article

[Rock lithological classification by hyperspectral, range 3D and color images](#)

Chemometrics and Intelligent Laboratory Systems, Volume 189, 2019, pp. 138-148

[Show abstract](#) ✓

Research article

[Sensory defects of virgin olive oil from a microbiological perspective](#)

Trends in Food Science & Technology, Volume 43, Issue 2, 2015, pp. 227-235

[Show abstract](#) ✓

Research article

[Phenolic compounds and quality parameters of family farming versus protected designation of origin \(PDO\) extra-virgin olive oils](#)

Journal of Food Composition and Analysis, Volume 43, 2015, pp. 75-81

[Show abstract](#) ✓

[View full text](#)

Copyright © 1999 Published by Elsevier B.V.



All content on this site: Copyright © 2023 Elsevier B.V., its licensors, and contributors. All rights are reserved, including those for text and data mining, AI training, and similar technologies. For all open access content, the Creative Commons licensing terms apply.



Meta 1018
Meta v. VL Collective
IPR2023-00923

ATTACHMENT 1B

Computer Networks: The International Journal of Computer and Telecommunications Net... 

very large conferences on the internet. the internet multimedia conferencing architecture

Authors:  [M. Handley](#),  [J. Crowcroft](#),  [C. Bormann](#),  [J. Ott](#) [Authors Info & Claims](#)

Computer Networks: The International Journal of Computer and Telecommunications Networking
, Volume 31, Issue 3 • Feb. 11, 1999 • pp 191–204 • [https://doi.org/10.1016/S0169-7552\(98\)00262-1](https://doi.org/10.1016/S0169-7552(98)00262-1)

Published: 11 February 1999 [Publication History](#)

 3  0

Feedback



Abstract

No abstract available.

Cited By

View all 

Shirmohammadi S, Ding L and Georganas N. (2019). An Approach for Recording Multimedia Collaborative Sessions. *Multimedia Tools and Applications*. **19**:2. (135-154). Online publication date: 1-Feb-2003.

<https://doi.org/10.1023/A:1022143111606>

CO

Meta 1018
Meta v. VL Collective
IPR2023-00923

<https://doi.org/10.1145/1008653.1008680>

Computer Networks: The International Journal of Computer and Telecommunications Net... ▼

(53-61).

<https://doi.org/10.1145/507670.507679>

Index Terms

Very large conferences on the Internet: the Internet multimedia conferencing architecture

Feedback

- > Computer systems organization
 - > Dependable and fault-tolerant systems and networks
 - > Embedded and cyber-physical systems
 - > Real-time systems
- > General and reference
 - ▼
Cross-computing tools and techniques
 - ▼
Performance
- > Human-centered computing



>

Information systems

Computer Networks: The International Journal of Computer and Telecommunications Net... ▼

▼

Multimedia information systems

>

World Wide Web

▼

Web applications

▼

Internet communications tools

▼

Web conferencing

Feedback

>

Networks

>

Network performance evaluation

>

Network protocols

▼

Transport protocols

>

Network types

▼

Public Internet



Comments

Computer Networks: The International Journal of Computer and Telecommunications Net... ▼

Feedback

View Issue's Table of Contents

Categories

Journals
Magazines
Books
Proceedings
SIGs
Conferences
Collections
People



About

About ACM Digital Library
ACM Digital Library Board
Subscription Information
Author Guidelines
Using ACM Digital Library
All Holdings within the ACM Digital Library
ACM Computing Classification System
Digital Library Accessibility

Meta 1018
Meta v. VL Collective
IPR2023-00923

Join ACM

Join SIGs

✉ Contact

f Facebook

Computer Networks: The International Journal of Computer and Telecommunications Net... ▼

Institutions and Libraries

i Feedback

i Bug Report

The ACM Digital Library is published by the Association for Computing Machinery. Copyright © 2023 ACM, Inc.

[Terms of Usage](#) | [Privacy Policy](#) | [Code of Ethics](#)



Feedback



ATTACHMENT 1C

MARC Bibliographic Record

LEADER 02624cas a2200625 a 4500

001 9943085403602122

005 20150409123053.0

008 990217c19999999ne mr p 0 a0eng c

010 \$asn 99047167

012 \$i9902\$m0

016 7_ \$a012403209\$2Uk

022 \$a1389-1286\$11389-1286

035 \$a(0CoLC)ocm40811573

035 \$9AXX5213UW

035 \$a4308540

035 \$a(WU)4308540-uwmadisondb

035 \$a(EXLNZ-01UWI_NETWORK)999852141902121

037 \$bElsevier Science, P.O. Box 945, New York, N.Y. 10159-0945\$c\$1249

040 \$aUHC\$beng\$cUHC\$dWAU\$dEYM\$dNYP\$dHUL\$dCUS\$dNSD\$dDLC\$dSYS\$dDLC\$dMYG\$dOCLCQ\$dNLGGC\$dLVB\$dOCLCQ\$dUKMGB\$dGZM

042 \$apcc

049 \$aGZMA

050 00 \$aTK5105.5\$b.C647

098 11 \$aS\$aC734\$aN388\$aIA1071

130 0_ \$aComputer networks (Amsterdam, Netherlands : 1999)

245 10 \$aComputer networks.

246 18 \$aComputer networks and ISDN systems\$f<Feb. 1999>-

260 \$aAmsterdam :\$bElsevier,\$c1999-

300 \$avolumes :\$billustrations ;\$c27 cm

310 \$aMonthly

336 \$atext\$btxt\$2rdacontent

337 \$aunmediated\$bn\$2rdamedia

338 \$avolume\$bn\$2rdacarrier

362 0_ \$aVol. 31, no. 1-2 (14 Jan. 1999)-

500 \$aTitle from cover.

500 \$a"The international journal of computer and telecommunications networking."

500 \$aSome issues have also distinctive titles.

530 \$aAvailable also online via the World Wide Web.

530 \$aTables of contents available via email through the ContentsDirect alerting service.

515 \$aSome issues published in combined form.

550 \$aOfficial journal of the International Council for Computer Communication.

650 _0 \$aComputer networks\$vPeriodicals.

650 _0 \$aIntegrated services digital networks\$vPeriodicals.

650 _0 \$aTelecommunication systems\$vPeriodicals.

650 _0 \$aDigital communications\$vPeriodicals.

650 17 \$aComputernetwerken.\$2gtt

650 _6 \$aRéseaux d'ordinateurs\$vPériodiques.

650 _6 \$aDonnées\$xTransmission\$vPériodiques.

650 _6 \$aTélécommunications, Systèmes de\$vPériodiques.

710 2_ \$aInternational Council for Computer Communication.

776 08 \$iOnline version:\$tComputer networks (Amsterdam, Netherlands : 1999 : Online)\$x1872-7069\$w(DLC) 2003242035\$w(OCOLC)42202506

776 1_ \$tComputer networks (Amsterdam, Netherlands : 1999)\$w(OCOLC)752418019

780 00 \$tComputer networks and ISDN systems\$x0169-7552\$w(DLC) 86641126\$w(OCOLC)12119244

850 \$aDLC\$aMiU\$aWaU

891 00 \$81\$av.\$bno.\$i(year)\$j(month)

891 40 \$81\$a<31>\$b<1/2>\$i<1999>\$j<Jan.>\$xprovisional

997 \$aMARCIVE

LEADER 02602nas--2200673-a-4500

001 991021976663702122

005 20221206191705.0

006 m-----o--d-----

007 cr-unu|||||

008 990820c19999999ne-mx-pso-----0---a0eng-c

010 \$a 2003242035

016 7_ \$a1499744 7\$2DE-600

016 7_ \$a012432534\$2Uk

016 7_ \$a012432534\$2Uk

016 7_ \$a012432534\$2Uk

016 7_ \$a012432534\$2Uk

016 7_ \$a012432534\$2Uk

022 \$a1872-7069

035 \$a(OCOLC)42202506

035 \$a(CKB)954925622232

035 \$a(CONSER)--2003242035

035 \$a(DE-599)ZDB1499744 7

035 \$a(EXLCZ)99954925622232

042 \$apcc

050 14 \$aTK5105.5\$b.C647

130 0_ \$aComputer networks (Amsterdam, Netherlands : 1999 : Online)

210 \$aCOMPUT. NETW

210 \$aCOMPUT NETW

210 \$aCOMPUTER NETWORKS THE INTERNATIONAL JOURNAL OF COMPUTER AND TELECOMMUNICATIONS NETWORKING

210 \$aCOMPUTER NETWORKS THE INTERNATIONAL JOURNAL OF COMPUTER AND TELECOMMUNICATIONS

210 \$aCOMPUT NETWORKS

222 _0 \$aComputer networks\$b(1999. Online)

245 10 \$aComputer networks\$h[e-journal].

246 \$aComputer networks

246 \$aComputer Networks & ISDN Systems

260 \$a[Amsterdam?] :\$bElsevier Science,\$c@1999-

310 \$a18 issues yearly,\$b2000-

321 \$aTwenty-four issues yearly,\$b1999

336 \$atext\$btxt\$2rdacontent

337 \$acomputer\$bc\$2rdamedia

338 \$aonline resource\$bcr\$2rdacarrier

362 0_ \$aVol. 31, issues 1/2 (14 Jan. 1999)-

500 \$aTitle from title information screen (viewed Mar. 26, 2003).

500 \$aSome issues have also a distinctive title.

550 \$aOfficial journal of the International Council for Computer Communication.

588 \$aLatest issue consulted: Vol. 41, issue 6 (22 Apr. 2003).

500 \$aRefereed/Peer-reviewed

650 _0 \$aComputer networks\$vPeriodicals.

650 _0 \$aIntegrated services digital networks\$vPeriodicals.

650 _0 \$aTelecommunication systems\$vPeriodicals.

650 _0 \$aDigital communications\$vPeriodicals.

650 _7 \$aComputer networks.\$2fast\$0(OCOLC)fst00872297

650 _7 \$aDigital communications.\$2fast\$0(OCOLC)fst00893634

650 _7 \$aIntegrated services digital networks.\$2fast\$0(OCOLC)fst00975649

650 _7 \$aTelecommunication systems.\$2fast\$0(OCOLC)fst01146038

655 _7 \$aPeriodicals.\$2fast\$0(OCOLC)fst01411641

710 2_ \$aInternational Council for Computer Communication.

780 00 \$tComputer networks and ISDN systems\$w(CKB)954926233417\$w(DLC)sf89091273\$x1879-2324

776 \$x1389-1286

906 \$aJOURNAL

MMS IDs**Document ID:**

999852141902121

Network Electronic IDs:

999852141902121, 9911025719702121

Network Physical IDs:

999852141902121

mms_mad_ids:

9943085403602122, 991021976663702122

mms_plt_ids:

991013837038802128, 991013840279002128, 991013884714702128

mms_lc_ids:

991016667320902125

mms_osh_ids:

991004595943802126

mms_sf_ids:

991015006401302130

mms_st_ids:

991013561168602131

mms_ww_ids:

991015831368802133

ATTACHMENT 1D

[The Library of Congress](#)>> [Go to Library of Congress Online Catalog](#)

LIBRARY OF CONGRESS AUTHORITIES

[Help](#)[New Search](#)[Search History](#)[Headings List](#)[Start Over](#)< [Previous](#)[Next](#) >[MARC Display](#)[Labelled Display](#)

LC control no.: sh2008101495

LCCN Permalink: <https://lcn.loc.gov/sh2008101495>

HEADING: Computer networks Periodicals

000 00448cz a2200169n 450

001 7438024

005 20110729160116.0

008 080206|| anannbabn |n ana

010 __ |a sh2008101495

035 __ |a (DLC)377031

035 __ |a (DLC)sh2008101495

040 __ |a DLC |b eng |c DLC

150 __ |a Computer networks |v Periodicals

667 __ |a Record generated for validation purposes.

670 __ |a Work cat.: Internet world,

906 __ |t 8888 |u te00 |v 0

953 __ |a te00

< [Previous](#) [Next](#) >**Save, Print or Email Records (View [Help](#))**Select Download Format: [Text \(Labelled Display\)](#)[Press to SAVE or PRINT](#)

Email Text (Labelled Display) to:

[Press to SEND EMAIL](#)[Help](#) - [Search](#) - [Search History](#) - [Headings List](#) - [Start Over](#)**Library of Congress**URL: <https://www.loc.gov/>

Mailing Address:

101 Independence Ave,
S.E.
Washington, DC 20540**Library of Congress Authorities**URL: <http://authorities.loc.gov/>**Library of Congress Online Catalog**URL: <https://catalog.loc.gov/>Questions, comments, error reports: [Contact Us](#)**Meta 1018**
Meta v. VL Collective
IPR2023-00923

ATTACHMENT 1E

[The Library of Congress](#)>> [Go to Library of Congress Online Catalog](#)

LIBRARY OF CONGRESS AUTHORITIES

[Help](#)[New Search](#)[Search History](#)[Headings List](#)[Start Over](#)< [Previous](#)[Next](#) >[MARC Display](#)[Labelled Display](#)

LC control no.: sh 86000852

LCCN Permalink: <https://lccn.loc.gov/sh86000852>

HEADING: Integrated services digital networks

000 00634cz a2200217n 450

001 4798998

005 20120328083840.0

008 860213i| anannbabn |a ana

010 __ |a sh 86000852

035 __ |a (DLC)sh 86000852

035 __ |a (DLC)145683

040 __ |a DLC |c DLC |d DLC

053 _0 |a TK5103.75

150 __ |a Integrated services digital networks

450 __ |a ISDN (Digital networks)

550 __ |w g |a Data transmission systems

550 __ |w g |a Digital communications

670 __ |a Work cat.: Tutorial : integrated services digital networks.

906 __ |t 8842 |u fk03 |v 0

952 __ |a 0 bib. record(s) to be changed

953 __ |a fg03 |b fg11

< [Previous](#)[Next](#) >**Save, Print or Email Records (View [Help](#))**Select Download Format: [Press to SAVE or PRINT](#)

Email Text (Labelled Display) to:

[Press to SEND EMAIL](#)[Help](#) - [Search](#) - [Search History](#) - [Headings List](#) - [Start Over](#)

Library of Congress

URL: <https://www.loc.gov/>

Library of Congress Authorities

URL: <http://authorities.loc.gov/>

Library of Congress Online Catalog

URL: <https://catalog.loc.gov/>

Meta v. VL Collective

IPR2023-00923



Mailing Address:
101 Independence Ave,

S.E.
Washington, DC 20540

Questions, comments, error reports: [Contact Us](#)

ATTACHMENT 1F

[The Library of Congress](#)>> [Go to Library of Congress Online Catalog](#)

LIBRARY OF CONGRESS AUTHORITIES

[Help](#)[New Search](#)[Search History](#)[Headings List](#)[Start Over](#)[Previous](#)[Next](#)[MARC Display](#)[Labelled Display](#)**LC control no.:** sh 85099890**LCCN Permalink:** <https://lccn.loc.gov/sh85099890>**HEADING:** Periodicals

000 01324cz a2200337n 450

001 4749694

005 20120717060125.0

008 860211 | anannbabn |b ana

010 __ |a sh 85099890

035 __ |a (DLC)sh 85099890

035 __ |a (DLC)6226910

035 __ |a (DLC)sp 85099890

035 __ |a (DLC)314952

040 __ |a DLC |c DLC |d DLC

053 _0 |a AP |c General periodicals

053 _0 |a PN4700 |b PN5650 |c History

150 __ |a Periodicals

360 __ |i subdivision |a Periodicals |i under specific subjects, e.g. |a Engineering--Periodicals; United States--History--Periodicals

450 __ |a Journals (Periodicals)

450 __ |a Magazines

550 __ |w g |a Library materials

550 __ |w g |a Mass media

550 __ |w g |a Serial publications

550 __ |a Newspapers

550 __ |a Press

670 __ |a UMI business vocab. |b (Journals)

680 __ |i Here are entered works on the periodicals of the world and, with appropriate subdivisions, works on special aspects or sections of periodicals. Works on periodicals in a specific language, or in a specific country or larger area, are entered under the adjectival form of the language or area, e.g. |a Arabic periodicals; Canadian periodicals.

906 __ |t 0420 |u te03 |v 0

910 __ |a Proposal saved by tc03 on 05/14/2012 at 17:15:30

953 __ |a xx00 |b td14

906 __ |t 1206 |v 0

Meta 1018
Meta v. VL Collective
IPR2023-00923

◀ Previous Next ▶

Save, Print or Email Records (View Help)	
Select Download Format:	Text (Labelled Display) ▼ Press to SAVE or PRINT
Email Text (Labelled Display) to:	Press to SEND EMAIL

[Help](#) - [Search](#) - [Search History](#) - [Headings List](#) - [Start Over](#)

**Library of Congress**URL: <https://www.loc.gov/>*Mailing Address:*101 Independence Ave,
S.E.
Washington, DC 20540**Library of Congress Authorities**URL: <http://authorities.loc.gov/>**Library of Congress Online Catalog**URL: <https://catalog.loc.gov/>**Questions, comments, error reports:** [Contact Us](#)

ATTACHMENT 1G

[The Library of Congress](#)>> [Go to Library of Congress Online Catalog](#)

LIBRARY OF CONGRESS AUTHORITIES

[Help](#)[New Search](#)[Search History](#)[Headings List](#)[Start Over](#)< [Previous](#)[Next](#) >[MARC Display](#)[Labelled Display](#)**LC control no.:** sh2010115961**LCCN Permalink:** <https://lcn.loc.gov/sh2010115961>**HEADING:** Telecommunication systems Periodicals

000 00485cz a2200169n 450

001 8410349

005 20110729164444.0

008 100914|| anannbabn |n ana

010 __ |a sh2010115961

035 __ |a (DLC)483399

035 __ |a (DLC)sh2010115961

040 __ |a DLC |b eng |c DLC

150 __ |a Telecommunication systems |v Periodicals

667 __ |a Record generated for validation purposes.

670 __ |a Work cat.: Computer networks and ISDN systems, [c1985]-

906 __ |t 8888 |u tc00 |v 0

953 __ |a tc00

< [Previous](#)[Next](#) >**Save, Print or Email Records (View [Help](#))**Select Download Format: [Text \(Labelled Display\)](#)[Press to SAVE or PRINT](#)

Email Text (Labelled Display) to:

[Press to SEND EMAIL](#)[Help](#) - [Search](#) - [Search History](#) - [Headings List](#) - [Start Over](#)**Library of Congress**URL: <https://www.loc.gov/>**Mailing Address:**101 Independence Ave,
S.E.
Washington, DC 20540**Library of Congress Authorities**URL: <http://authorities.loc.gov/>**Library of Congress Online Catalog**URL: <https://catalog.loc.gov/>**Questions, comments, error reports:** [Contact Us](#)**Meta 1018
Meta v. VL Collective
IPR2023-00923**

ATTACHMENT 1H

The Library of Congress

>> [Go to Library of Congress Online Catalog](#)

LIBRARY OF CONGRESS AUTHORITIES



Help

New Search

Search History

Headings List

Start Over

< Previous

Next >

MARC Display

Labelled Display

LC control no.: sh 85037972

LCCN Permalink: <https://lccn.loc.gov/sh85037972>

HEADING: Digital communications

000 00785cz a2200277n 450

001 4690005

005 20120323072233.0

008 860211i| anannbabn |b ana

010 __ |a sh 85037972

035 __ |a (DLC)sh 85037972

035 __ |a (DLC)36690

040 __ |a DLC |c DLC |d DLC

053 _0 |a TK5103.7 |b TK5103.75

150 __ |a Digital communications

450 __ |a Communications, Digital

450 __ |a Digital transmission

450 __ |a Pulse communication

550 __ |w g |a Digital electronics

550 __ |w g |a Pulse techniques (Electronics)

550 __ |w g |a Telecommunication

550 __ |a Digital media

550 __ |a Signal processing |x Digital techniques

670 __ |a UMI business vocab. |b (Digital transmission)

906 __ |t 9843 |u te07 |v 1

952 __ |a LC-BRF

953 __ |a xx00 |b ta25

< Previous

Next >

Save, Print or Email Records (View [Help](#))

Select Download Format: Text (Labelled Display) v

Press to SAVE or PRINT

Email Text (Labelled Display) to:

Press to SEND EMAIL

Meta 1018
Meta v. VL Collective
IPR2023-00923

[Help](#) - [Search](#) - [Search History](#) - [Headings List](#) - [Start Over](#)

**Library of Congress**

URL: <https://www.loc.gov/>

Mailing Address:

101 Independence Ave,
S.E.
Washington, DC 20540

Library of Congress Authorities

URL: <http://authorities.loc.gov/>

Library of Congress Online Catalog

URL: <https://catalog.loc.gov/>

Questions, comments, error reports: [Contact Us](#)

ATTACHMENT 1I

OCLC 40811573 No holdings in XXX - 185 other holdings

Re		En 199		Re 2019032		
c		tere 902		place 7011329		
Stat		c 17		d .7		
Type		ELvl	Srce	GPub	Ctrl	Lang
BLvl		Form	Conf	Freq	MRec	Ctry
S/L		Orig	EntW	Reql	Alph	
Desc		SrTp	Cont	DtSt	Dates	

Continuing Resources

007 c \$b r \$d u \$e n \$f u

010 sn 99047167

040 UHC \$b eng \$c UHC \$d WAU \$d EYM \$d NYP \$d HUL \$d CUS \$d NSD \$d DLC \$d SYS \$d DLC \$d MYG
\$d OCLCQ \$d NLGGC \$d LVB \$d OCLCQ \$d UKMGB \$d OCLCF \$d OCLCO \$d I8M \$d OCLCQ \$d GILDS \$d
OCLCO \$d ZAS \$d OCLCO

012 \$i 9902 \$m 0

015 GB9916147 \$2 bnb

016 7 012403209 \$2 Uk

019 263091800 \$a 860141235 \$a 1013282969 \$a 1027260450 \$a 1064547366

022 1389-1286 \$l 1389-1286

037 \$b Elsevier Science, P.O. Box 945, New York, N.Y. 10159-0945 \$c \$\$1249

042 pcc

050 00 TK5105.5 \$b .C647

090 \$b

130 0 Computer networks (Amsterdam, Netherlands : 1999)

245 10 Computer networks.

246 18 Computer networks and ISDN systems \$f <Feb. 1999-

260 Amsterdam : \$b Elsevier, \$c 1999-

300 volumes : \$b illustrations ; \$c 27 cm

310 Monthly

336 text \$b txt \$2 rdacontent

337 unmediated \$b n \$2 rdamedia

338 volume \$b nc \$2 rdacarrier

362 0 Vol. 31, no. 1-2 (14 Jan. 1999)-

500 Title from cover.

500 Some issues have also distinctive titles.

515 Some issues published in combined form.

530 Also issued online via the World Wide Web.

530 Tables of contents available via email through the ContentsDirect alerting service.

550 Official journal of the International Council for Computer Communication.

588 Latest issue consulted: Vol. 32, no. 3 (Mar. 2000).

650 0 Computer networks \$v Periodicals.

650 0 Integrated services digital networks \$v Periodicals.

650 0 Telecommunication systems \$v Periodicals.

650 0 Digital communications \$v Periodicals.

650 6 Réseaux d'ordinateurs \$v Périodiques.

650 6 Données \$x Transmission \$v Périodiques.

650 6 Télécommunications, Systèmes de \$v Périodiques.
 650 7 Computer networks. \$2 fast \$0 (OCoLC)fst00872297
 650 7 Digital communications. \$2 fast \$0 (OCoLC)fst00893634
 650 7 Integrated services digital networks. \$2 fast \$0 (OCoLC)fst00975649
 650 7 Telecommunication systems. \$2 fast \$0 (OCoLC)fst01146038
 650 17 Computernetwerken. \$2 gtt
 655 7 Periodicals. \$2 fast \$0 (OCoLC)fst01411641
 655 7 Periodicals. \$2 lcgft
 710 2 International Council for Computer Communication.
 776 1 \$t Computer networks (Amsterdam, Netherlands : 1999 : Online) \$w (DLC) 2003242035 \$w
 (OCoLC)42202506
 776 08 \$i Online version: \$t Computer networks (Amsterdam, Netherlands : 1999) \$w
 (OCoLC)752418019
 780 00 \$t Computer networks and ISDN systems \$x 0169-7552 \$w (DLC) 86641126 \$w
 (OCoLC)12119244
 850 DLC \$a MiU \$a WaU
 856 41 \$u <http://www.sciencedirect.com/science/journal/13891286>

ATTACHMENT 1J

S
C734
N388
IA1071

COMPUTER NETWORKS



**The International Journal of Computer
and Telecommunications Networking**

Special Issue

Internet Telephony

Guest Editors: M. Graf and H.L. Truong

NOW included in your subscription:

**ELECTRONIC
ACCESS**

www.elsevier.nl/locate/elecacc

COMPUTER NETWORKS EDITORIAL BOARD

The International Journal of Computer and Telecommunications Networking

EDITOR-IN-CHIEF:

Philip H. Enslow, Jr.
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280, USA
Tel.: +1 404 894 3187
Fax: +1 404 894 3188
Email: enslow@cc.gatech.edu
MCIMail: PEnslow

ASSOCIATE EDITOR-IN-CHIEF:

Harry Rudin
Vordere Bergstrasse 1
CH-8942 Oberrieden
Switzerland
Tel.: +41-1-7203707
Fax: +41-1-7203777
Email: hrudin@smile.ch

Sudhir Aggarwal
Dept. of Computer Science
State University of New York
P.O. Box 6000
Binghamton, NY 13902-6000
USA

I.F. Akyildiz
School of Electrical Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0250
USA
Email: ian@ee.gatech.edu

P.D. Amer
103 Smith Hall
Dept. Comp and Inf. Sc.
Univ. of Delaware
Newark, DE 19716
USA

M. Ammar
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
USA
Email: ammar@cc.gatech.edu

Harmen R. van As
Institute of Communications Networks
Vienna University of Technology
Gusshausstrasse 25/388
A-1040 Vienna, Austria

Gregor v. Bochmann
University of Montreal
Dept. of Computer Science
and Operational Research
P.O. Box 6128, Station A
Montreal, Quebec H3C 3J7
Canada
Email: bochmann@iro.umontreal.ca

L.G. Bouma
KPN Research
P.O. Box 421
2260 AK Leidschendam
Netherlands
E-mail: l.g.bouma@research.kpn.com

V.G. Cerf, VP
MCI Telecommunications Corporation
2100 Reston Parkway
Reston, VA 22091
USA
Email: vcferf@isoc.org

Imrich Chiantac
Electrical, Computer & Systems Engrg.
Boston University
44 Cunningham Street
Boston, MA 02215
USA
Email: chiantac@bu.edu

Magda Elzarki
Dept. of Electrical Engineering
University of Pennsylvania
Philadelphia, PA 19104-6390
USA
Email: magda@ee.upenn.edu

Janusz Filipiak
Telecommunications Dept.
The University of Mining and Metallurgy
Al. Mickiewicza 30
30-059 Krakow, Poland
Email: gefilipi@cyf-kr.edu.pl

N.L.S. Fonseca
Institute of Computing
State University of Campinas
P.O. Box 6176
13083-970 Campinas SP
Brazil
Email: nfonseca@dcc.unicamp.br

Luigi Fratta
Dipart. di Elettronica e Informazione
Politecnico di Milano
Piazza L. da Vinci 32
I-20133 Milano, Italy
Email: fratta@ipmle2.elet.polimi.it

Aura Ganz
Dept. of Elec. and Comp. Eng.
University of Massachusetts
Amherst, MA 01003, USA
Email: ganz@risky.ecs.umass.edu

N.D. Georganas
Multimedia Communications Research
Lab. (MCRLab)
Dept. of Electrical Engrg.
University of Ottawa
161 Louis Pasteur Priv., Room A-617
P.O. Box 450, Station "A"
Ottawa, Ontario
Canada K1N 6N5
Email: georgana@mcrlab.uottawa.ca

Z.J. Haas
School of Electrical Engrg.
Wireless Networks Laboratory
Cornell University
323 Frank Rhodes Hall
Ithaca, NY 14853
USA
Email: haas@ee.cornell.edu

Raj Jain
Ohio State University
Department of Computer
and Information Science
2015 Neil Avenue Mall
Room 497, Dreese Labs
Columbus, OH 43210-1277
USA

Arik N. Kashper
Technical Manager
AT&T Bell Labs, Room 2J-407
101 Crawford Corner Road
Holmdel, NJ 07733, USA
Email: arik@attmail.com

Leonard Kleinrock
University of California
Computer Science Dept.
3732 Boelter Hall
Los Angeles, CA 90024
USA
Email: lk@cs.ucla.edu

Chin-Tau Lea
School of Electrical Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0250
USA
Email: leact@ee.gatech.edu

Luciano Lenzini
University of Pisa
Department of Information Engineering
Via Diotisalvi 2
56126 Pisa, Italy

Amar Lindner (Patent Editor)
European Patent Office
Patentlaan 2
P.O. Box 5818
2280 HV Rijswijk
Netherlands
Email: lindner@epo.e-mail.com

Eric Manning
Faculty of Engineering
University of Victoria
P.O. Box 3055, MS 7209
Victoria, BC
Canada V8W 3P5
Email: Eric.Manning@enr.UVic.ca

Bjorn Pehrson
KTH/SICS
Electrum 204
S164 40 Stockholm-Kista
Sweden
Email: bjorn@it.kth.se

Ramon Puigjaner
Universitat de les Illes Balears
Dept. de c. Matemàtiques
i Informàtica
Ctra. de Valldemossa KM 7.6
07071 Palma
Spain
Email: dmirp0@ps.vib.es

Guy Pujolle
Head, MASI Lab
University of Versailles
45 Av. des Etats-Unis
78035 Versailles Cedex
France
Email: pujolle@masi.ibp.fr

Juan Quemada
ETSI Telecomunicacion
Ciudad Universitaria S/N
28040 Madrid
Spain
Email: jquemada@dct.upm.es

S.V. Raghavan
Dept. of Computer Science
and Engineering
Indian Institute of Technology
600036 Madras
India
Email: svr@iitm.ernet.in

James W. Roberts
France Telecom
CNET/PAA/KTR
38 rue du Général Leclerc
92131 Issy les Mulinéaux
France
Email: roberts@issy.cnet.fr

Nobuhiko Shimasaki
Tokai University
Institute of Research and
Development
2-28-4 Tomigaya, Shibuya-ku
Tokyo 151
Japan
Email: shimasak@keyaki.cc.u-tokai.ac.jp

Kazem Sohraby
AT&T Bell Labs, Room 3M-323
101 Crawford Corner Road
Holmdel, NJ 07733
USA
Email: sohraby@att.com

Khosrow Sohraby
Comput. Sci. Telecomm. Program
University of Missouri
5100 Rockhill Road
Kansas City, MO 64110-2499
USA
Email: sohraby@cstp.umkc.edu

Csaba Szabo
Technical University of Budapest
Dept. of Telecommunications
Sztocek u.2.
H-1111 Budapest
Hungary
Email: szabo@com.hit.bme.hu

Mehmet Ulema
Wireless Communications
Technologies
Daewoo Telecom, Ltd.
New Jersey R&D Center
1715 Highway 35 North, Suite 302
Middletown, NJ 07748
USA
Email: ulema@daewootelecom.com

Kenneth S. Vastola
Electrical, Computer and Systems
Engineering Dept.
JEC-6024
Rensselaer Polytechnic Institute
110 8th Street
Troy, NY 12180-3590
USA
Email: vastola@rpi.edu



COMPUTER NETWORKS is the official Journal of the International Council for Computer Communication, P.O. Box 9745, Washington, D.C. 20016, U.S.A.

President: *Pramode K. Verma*
Executive Vice-President: *S. Ramani*
Secretary General: *Masao Kato*
Deputy Secretary General: *Karen Duncan*
Executive Secretary: *David Balson*
Treasurer: *John McKendree*
Vice-President: *Trevor Housley*
Vice-President: *Jehuda Kella*
Vice-President: *Dipak Khakhar*
Vice-President: *Michael Miller*

The Council is organized to foster scientific research, development, and applications of computer communication; to foster progress in evaluation of applications of computer communication; to foster study of the potential social and economic impacts of computer communication and of policies which may strongly influence those impacts; to improve public understanding of computer communication and its application; and to publish and disseminate printed matter, as appropriate.

COMPUTER NETWORKS features a regular news column of "COST-11 Ter".

Meta 1018
Meta v. VL Collective
IPR2023-00923

LIBRARY RECOMMENDATION CARD

Use this card to recommend that your library subscribes to **COMPUTER NETWORKS**

ROUTE VIA INTERDEPARTMENTAL MAIL

To: The Serials Librarian at _____
From: _____ Title: _____
Dept: _____ Phone: _____

I recommend that our library takes a subscription to **COMPUTER NETWORKS** (ISSN 1389-1286). Please note that a subscription to **COMPUTER NETWORKS** gives all library users free electronic access to Full-Text Articles (in PDF format), International Industry News, Tables of Contents, Abstracts, Free-Text Search Facilities and more. For further details: www.elsevier.nl/locate/comnet

The Subscription Price for 1999, Volume 31 (24 issues) is NLG 2434.00 / US\$ 1235.00 including postage and handling.

Signature _____ Date _____

ORDER CARD

☐ Please enter a 1999 subscription to **COMPUTER NETWORKS**, Volume 31 (24 issues)
ISSN 1389-1286 (formerly ISSN 0169-7552). Subscription price: NLG 2434.00 / US\$ 1235.00

Please print clearly

Name _____ Initials _____
Position _____ Title _____
Institute/Company _____
Address _____
City _____ State _____
Zip/Postal Code _____ Country _____
Tel _____ Fax _____
E-mail _____ VAT number _____

0501069A

Payment Details

- ☐ Please send me a pro forma invoice
☐ Payment enclosed (*make payable to Elsevier Science*):
☐ Cheque ☐ UNESCO coupons
☐ I wish to pay by credit card (*accepted from individuals only*). Your credit card will be debited including VAT when applicable.
☐ American Express ☐ Eurocard/MasterCard/Access ☐ VISA
Card number _____ Exp. Date _____
Signature _____ Date _____

Dutch Guilder (NLG) prices apply to customers within Europe and Japan. US Dollar (US\$) prices apply to customers in all other countries.

Meta 1018
Meta v. VL Collective
IPR2023-00923

Send to your regular supplier or to the nearest address overleaf

Addresses: *New York*
Elsevier Science
Customer Support Department
P.O. Box 945
New York, NY 10159-0945
USA
TEL: (+1) 212-633-3730
Toll-free for customers in the USA and Canada: 1-888-437-4636
FAX: (+1) 212-633-3680
E-MAIL: usinfo-f@elsevier.com

Amsterdam
Elsevier Science
Customer Support Department
P.O. Box 211
1000 AE Amsterdam
The Netherlands
TEL: (+31) 20-485-3757
FAX: (+31) 20-485-3432
E-MAIL: nlinfo-f@elsevier.nl

NORTH-HOLLAND **Meta 1018**
An Imprint of Elsevier Science **Meta y. VL Collective**
IPR2023-00923



COMPUTER NETWORKS

The International Journal of Computer
and Telecommunications Networking

KURT F. WENDT LIBRARY
COLLEGE OF ENGINEERING

MAR - 1 1999

UW-MADISON, WI 53706

Volume 31, Number 3, 11 February 1999

Internet Telephony

Guest Editors: M. Graf and H.L. Truong



ELSEVIER Amsterdam – Lausanne – New York – Oxford – Shannon – Tokyo

Meta 1018
Meta v. VL Collective
IPR2023-00923

This journal and the individual contributions contained in it are protected by the copyright of Elsevier Science B.V., and the following terms and conditions apply to their use:

Photocopying

Single photocopies of single articles may be made for personal use as allowed by national copyright laws. Permission of the publisher and payment of a fee is required for all other photocopying, including multiple or systematic copying, copying for advertising or promotional purposes, resale, and all forms of document delivery. Special rates are available for educational institutions that wish to make photocopies for non-profit educational classroom use.

Permissions may be sought directly from Elsevier Science Rights & Permissions Department, PO Box 800, Oxford OX5 1DX, UK; phone: (+44) 1865 843830, fax: (+44) 1865 853333, e-mail: permissions@elsevier.co.uk. You may also contact Rights & Permissions directly through Elsevier's home page (<http://www.elsevier.nl>), selecting first 'Customer Support', then 'General Information', then 'Permissions Query Form'.

In the USA, users may clear permissions and make payments through the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA; phone: (978) 7508400, fax: (978) 7504744, and in the UK through the Copyright Licensing Agency Rapid Clearance Service (CLARCS), 90 Tottenham Court Road, London W1P 0LP, UK; phone: (+44) 171 436 5931; fax: (+44) 171 436 3986. Other countries may have a local reprographic rights agency for payments.

Derivative Works

Subscribers may reproduce tables of contents or prepare lists of articles including abstracts for internal circulation within their institutions. Permission of the publisher is required for resale or distribution outside the institution.

Permission of the publisher is required for all other derivative works, including compilations and translations.

Electronic Storage or Usage

Permission of the publisher is required to store or use electronically any material contained in this journal, including any article or part of an article. Contact the publisher at the address indicated.

Except as outlined above, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the publisher.

Address permissions requests to: Elsevier Science Rights & Permissions Department, at the mail, fax and e-mail addresses noted above.

Notice

No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made.

Although all advertising material is expected to conform to ethical (medical) standards, inclusion in this publication does not constitute a guarantee or endorsement of the quality or value of such product or of the claims made of it by its manufacturer.

Contents

<i>M. Graf and H.L. Truong</i> Internet Telephony	153
<i>J. Rinde</i> Telephony in the year 2005	157
<i>R. Guérin and V. Peris</i> Quality-of-service in packet networks: basic mechanisms and directions	169
<i>M. Handley, J. Crowcroft, C. Bormann and J. Ott</i> Very large conferences on the Internet: the Internet multimedia conferencing architecture	191
<i>J. Toga and J. Ott</i> ITU-T standardization activities for interactive multimedia communications on packet-based networks: H.323 and related recommendations	205
<i>M. Reid</i> Multimedia conferencing over ISDN and IP networks using ITU-T H-series recommendations: architecture, control and coordination	225
<i>H. Schulzrinne and J. Rosenberg</i> Internet Telephony: architecture and protocols – an IETF perspective	237
<i>J.-P. Hubaux, C. Gbaguidi, S. Koppenhoefer and J.-Y. Le Boudec</i> The impact of the Internet on telecommunication architectures	257
Calendar	275



Abstracted/Indexed in: ABI Inform, ACM Computing Reviews, Cambridge Scientific Abstracts, Compumath Citation Index, Computer Abstracts, Computer Literature Index, Current Contents: Engineering, Computing & Technology, Engineering Index/Compendex, INSPEC Information Services, Library and Information Science Abstracts (LISA), Research Alert, Scisearch, Sociological Abstracts, Zentralblatt für Mathematik.

Special Issue

Internet Telephony

Guest Editors: M. Graf and H.L. Truong



This journal is part of **ContentsDirect**, the *free* alerting service which sends tables of contents by e-mail for Elsevier Science books and journals. The quickest way to register for **ContentsDirect** is via the World Wide Web at: www.elsevier.nl/locate/ContentsDirect

If you don't have access to the WWW you can register for this service by sending an e-mail message to cdsubs@elsevier.co.uk specifying the title of the publication you wish to register for. The tables of contents are also available on the Elsevier Science website at: www.elsevier.nl or www.elsevier.com or www.elsevier.co.jp



S
C734
N388
IA1071

COMPUTER NETWORKS



**The International Journal of Computer
and Telecommunications Networking**

Special Issue

Internet Telephony

Guest Editors: M. Graf and H.L. Truong

NOW included in your subscription:

**ELECTRONIC
ACCESS**

www.elsevier.nl/locate/elecacc

Meta 1018

Meta v. VL Collective

IPR2023-00923

This journal and the individual contributions contained in it are protected by the copyright of Elsevier Science B.V., and the following terms and conditions apply to their use:

Photocopying

Single photocopies of single articles may be made for personal use as allowed by national copyright laws. Permission of the publisher and payment of a fee is required for all other photocopying, including multiple or systematic copying, copying for advertising or promotional purposes, resale, and all forms of document delivery. Special rates are available for educational institutions that wish to make photocopies for non-profit educational classroom use.

Permissions may be sought directly from Elsevier Science Rights & Permissions Department, PO Box 800, Oxford OX5 1DX, UK; phone: (+44) 1865 843830, fax: (+44) 1865 853333, e-mail: permissions@elsevier.co.uk. You may also contact Rights & Permissions directly through Elsevier's home page (<http://www.elsevier.nl>), selecting first 'Customer Support', then 'General Information', then 'Permissions Query Form'.

In the USA, users may clear permissions and make payments through the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, USA; phone: (978) 7508400, fax: (978) 7504744, and in the UK through the Copyright Licensing Agency Rapid Clearance Service (CLARCS), 90 Tottenham Court Road, London W1P 0LP, UK; phone: (+44) 171 436 5931; fax: (+44) 171 436 3986. Other countries may have a local reprographic rights agency for payments.

Derivative Works

Subscribers may reproduce tables of contents or prepare lists of articles including abstracts for internal circulation within their institutions. Permission of the publisher is required for resale or distribution outside the institution.

Permission of the publisher is required for all other derivative works, including compilations and translations.

Electronic Storage or Usage

Permission of the publisher is required to store or use electronically any material contained in this journal, including any article or part of an article. Contact the publisher at the address indicated.

Except as outlined above, no part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without prior written permission of the publisher.

Address permissions requests to: Elsevier Science Rights & Permissions Department, at the mail, fax and e-mail addresses noted above.

Notice

No responsibility is assumed by the Publisher for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions or ideas contained in the material herein. Because of rapid advances in the medical sciences, in particular, independent verification of diagnoses and drug dosages should be made.

Although all advertising material is expected to conform to ethical (medical) standards, inclusion in this publication does not constitute a guarantee or endorsement of the quality or value of such product or of the claims made of it by its manufacturer.

Very large conferences on the Internet: the Internet multimedia conferencing architecture

M. Handley ^{a,*}, J. Crowcroft ^b, C. Bormann ^c, J. Ott ^c

^a *Information Sciences Institute, University of Southern California, Marina Del Rey, USA*

^b *Department of Computer Science, University College London, London, UK*

^c *Informatics Department, University of Bremen, Bremen, Germany*

Abstract

In this paper we provide an overview of multimedia conferencing on the Internet. The protocols mentioned are all specified elsewhere as Internet-Drafts or RFCs. Each RFC gives details of the protocol itself, how it works and what it does. This document attempts to provide the reader with an overview of how the components fit together and of some of the assumptions made, as well as some statement of direction for those components still in a nascent stage. © 1999 Published by Elsevier Science B.V. All rights reserved.

Keywords: Multimedia; Conferencing; Multicast; Quality of service; Real-time transport

1. Introduction

The Internet is not currently very good at carrying audio and video. This is hardly surprising as it was not designed or engineered with real-time traffic in mind, but there has recently been a great deal of interest in using the Internet for telephony services. Part of this has come from pricing anomalies that make Internet telephony somewhat artificially cheaper than traditional telephone services, but this is not the whole story. The Internet itself is improving to better handle traffic such as audio and video, and in the medium term, the Internet should be able to provide good quality realtime multimedia services, although such quality improvements are likely to incur additional charges.

However, the real interest in using the Internet for audio and video should come from the prospect for a

single ubiquitous communications network that not only allows traditional telephony services, but also video, shared collaboration tools, and through IP Multicast, multi-party conferences and multimedia sessions that scale from small group meetings through to television-sized audiences. In principle, this may lead to a “democratization” of telecommunication services, where licenses to broadcast are not required to control physical access to the limited broadcast medium (although they may still be required for political reasons).

It is far from clear what services will eventually emerge using such communications capabilities. We can only say that the technical capability to have large numbers of sessions ranging in audience from hundreds to millions of participants, largely unlimited by geographic boundaries, will lead to services and social structures that do not exist today. However, we can describe the basic technologies that are likely to bring about such changes, and in this paper

* Corresponding author. E-mail: mjh@east.isi.edu.

we attempt to provide such an overview. We leave it to the reader to imagine the uses to which this technology will be put.

2. The technology

In conjunction with computers, the term “conferencing” is often used in two different ways: firstly, to refer to the *asynchronous* exchanges of messages between multiple users, as with bulletin boards or mailing lists; secondly, to refer to *synchronous* or so-called “real-time” conferencing, including audio, video, shared whiteboards and other applications. This paper is about the architecture for this latter application in an Internet environment.

There are other infrastructures for conferencing in the world: POTS (Plain Old Telephone System) networks often provide voice conferencing and phone-bridges, while ISDN provides H.320 [16] for small, strictly organised video-telephony conferencing. The architecture that has evolved in the Internet is far more general as well as being scalable to very large groups, and permits the open introduction of new media and new applications as they are devised.

The key factors of a conferencing architecture are how it supports communication between possibly large numbers of people and the real-time delivery of information. In the Internet, this is supported at a number of levels in the protocol stack. The remainder of this section provides an overview of these features, and the rest of the document describes each aspect in more detail.

In a conference, information is typically distributed to all the conference participants. Early conferencing systems used a fan-out of data streams, e.g., one connection between each pair of participants, which meant that the same information must cross some networks more than once. The Internet architecture uses the more efficient approach of *multicasting* the information to all participants (Section 3).

Multimedia conferences require real-time delivery of at least the audio and video information streams used in the conference. In an ISDN context, fixed rate circuits are allocated for this purpose—whether their bandwidth is required at any particular instance

or not. On the other hand, the traditional Internet service model (“best effort”) cannot make the necessary quality of service available in congested networks. New service models are being defined in the Internet together with protocols to *reserve* capacity or *prioritise* traffic in a more flexible way than that available with circuit switching (Section 4). The nature of the Internet reflects that of the world in that it is very heterogeneous. Techniques exist to exploit this, and to deliver appropriate quality to different participants in the same conference according to their capabilities.

In a datagram network, multimedia information must be transmitted in packets, some of which may be delayed more than others. For audio and video streams to be played out at the recipient with the correct timing, information must be transmitted that allows the recipient to reconstitute this timing. A *transport protocol* with the specific functions needed for this has been defined (Section 5).

The people participating in a conference generally need to have a specific idea of the context in which the conference is happening, which can be formalized as a conference *policy*. Some conferences are essentially crowds gathered around an attraction, while others have very formal guidelines on who may take part (listen in) and who may speak at which point. In any case, initially the participants must find each other, i.e. establish communication relationships (conference *setup*, Section 7). During the conference, some conference *control* information is exchanged to implement a conference policy or at least to inform the crowd of who is present (Section 6).

In addition, *security* measures may be required to actually enforce the conference policy; for example, to control who is listening and to authenticate contributions as actually originating from a specific person. In the Internet, there is little tendency to rely on the traditional “security” of distribution offered by the phone system. Instead, cryptographic methods are used for encryption and authentication, which need to be supported by additional conference setup and control mechanisms (Section 8).

The protocol stacks that support Internet multimedia conferencing are shown in Fig. 1. Most of the protocols are not deeply layered unlike many protocol stacks, but rather are used in a complementary

Meta 1018

Meta v. VL Collective

IPR2023-00923




Fig. 1. Internet multimedia conferencing protocol stacks.

fashion for specific tasks to produce a complete conference.

3. Multicast traffic distribution

IP multicast provides efficient many-to-many data distribution in an Internet environment. It is easy to view IP multicast as simply an optimisation for data distribution, and indeed this is the case, but IP multicast can also result in a different way of thinking about application design. To see why this might be the case, examine the IP multicast service model, as described by Van Jacobson [9]:

- Senders just send to *the group*.
- Receivers express an interest in receiving data sent to *the group*.
- Routers conspire to deliver data from senders to receivers.

With IP multicast, the group is indirectly identified by a single IP class-D multicast address.

Several things are important about this service model from an architectural point of view. Receivers do not need to know who or where the senders are to receive traffic from them. Senders never *need* to know who the receivers are. Neither senders or receivers need care about the network topology as the network optimises delivery.

An IP multicast group is scalable because information about group membership and group changes at the IP level are kept local to routers near the relevant members. How this is performed depends on the particular multicast routing scheme in use local to the member, and although it is not a trivial task, several solutions do exist and therefore multicast routing will not be discussed in detail here. For more

detailed information on multicast routing, see [1,5,7,8,22]. Typically, as a group with s senders and r receivers increases in size, state in routers scales $O(s)$ or $O(1)$ depending on the routing scheme in use. This state may be in on-tree routers for newer so-called sparse-mode algorithms such as PIM, or in off-tree routers for older so-called dense-mode algorithms such as DVMRP. Thus the most scalable current multicast routing algorithms require $O(1)$ state in on-tree routers, and hence the total routing state scales $O(g)$ in a router that is on-tree for g groups. We can also envisage multicast routing schemes which require less than $O(g)$ state¹, but the requirement is not currently urgent, so none of these have yet been implemented.

The level of indirection introduced by the IP class D address denominating the group solves the distributed systems binding problem, by pushing this task down into routing; given a multicast address (and UDP port), a host can send a message to the members of a group without needing to discover who they are. Similarly receivers can “tune in” to multicast data sources without needing to bother the data source itself with any form of request.

IP multicast is a natural solution for multi-party conferencing because of the efficiency of the data distribution trees, with data being replicated in the network at appropriate points rather than in end-systems. It also avoids the need to configure special-purpose servers to support the session, which require support, and which cause traffic concentration and can be a bottleneck. For larger broadcast-style sessions, it is essential that data-replication be carried out in a way that only requires per-receiver network-state to be local to each receiver, and that data-replication occurs within the network. Attempting to configure a tree of application-specific replication servers for such broadcasts rapidly becomes a “multicast routing” problem, and thus native multicast support is a more appropriate solution.

¹ With IP encapsulation, not all on-tree routers need hold the state for a group whose traffic they are forwarding—traffic for the group can be encapsulated (either unicast or multicast) between on-tree routers nearer the edge of the network, reducing some of the state burden on backbone routers.

3.1. Address allocation

How does an application choose a multicast address to use?

In the absence of any other information, we can bootstrap a multicast application by using *well-known* multicast addresses. Routing (unicast and multicast) and group membership protocols [6] can do just that. However, this is not the best way of managing applications of which there is more than one instance at any one time.

For these, we need a mechanism for allocating group addresses dynamically, and a directory service which can hold these allocations together with some key (session information for example—see later), so that users can look up the address associated with the application. The address allocation and directory functions should be distributed to scale well.

Multicast address allocation is currently an active area of research. For many years multicast address allocation has been performed using multicast session directories (Section 7.1), but as the users and uses of IP multicast increase, it is becoming clear that a more hierarchical approach is required.

An architecture [13] is currently being developed based around a well-defined API that an application can use to request an address. The host then requests an address from a local address allocation server, which in turn chooses and reserves an unallocated address from a range dynamically allocated to the domain. By allocating addresses in a hierarchical and topologically sensitive fashion, the address itself can be used in a hierarchical multicast routing protocol currently being developed (BGMP [29]) that will help multicast routing scale more gracefully than current schemes.

4. Internet service models

Traditionally the Internet has provided so-called *best-effort* delivery of datagram traffic from senders to receivers. No guarantees are made regarding when or if a datagram will be delivered to a receiver, however datagrams are normally only dropped when a router exceeds a queue size limit due to congestion. The best-effort Internet service model does not as-

sume FIFO queuing, although many routers have implemented this.

With best-effort service, if a link is not congested, queues will not build at routers, datagrams will not be discarded in routers, and delays will consist of serialisation delays at each hop plus propagation delays. With sufficiently fast link speeds, serialisation delays are insignificant compared to propagation delays².

If a link is congested, with best-effort service, queuing delays will start to influence end-to-end delays, and packets will start to be lost as queue size limits are exceeded. Real-time traffic does not cope terribly well with packet loss levels of more than a few percent, although it is possible to add redundancy [14] to increase the levels at which loss becomes a problem. In the last few years a significant amount of work has also gone into providing non-best-effort services that would provide a better assurance that an acceptable quality conference will be possible.

4.1. Non-best effort service

Real-time Internet traffic is defined as datagrams that are delay sensitive. It could be argued that all datagrams are delay sensitive to some extent, but for these purposes we refer only to datagrams where exceeding an end-to-end delay bound of a few hundred milliseconds renders the datagrams useless for the purpose they were intended. For the purposes of this definition, TCP traffic is normally not considered to be real-time traffic, although there may be exceptions to this rule.

On congested links, best-effort service queuing delays will adversely affect real-time traffic. This does not mean that best-effort service cannot support real-time traffic—merely that congested best-effort links seriously degrade the service provided. For such congested links, a “better-than-best-effort” service is desirable.

To achieve this, the service model of the routers can be modified. FIFO queuing can be replaced by packet forwarding strategies that discriminate differ-

² For slow links, a set of mechanisms has been defined that helps minimize serialisation and link access delays [2,18].

ent “flows” of traffic. The idea of a flow is very general. A flow might consist of “all marketing site web traffic”, or “all file server traffic to and from teller machines”. On the other hand, a flow might consist of a particular sequence of packets from an application in a particular machine to a peer application in another particular machine set up on request, or it might consist of all packets marked with a particular Type-of-Service bit.

There is really a spectrum of possibilities for non-best-effort service something like that shown in Fig. 2. This spectrum is intended to illustrate that between best-effort, and hard per-flow guarantees lie many possibilities for non-best-effort service, including having hard guarantees based on an aggregate reservation, assurances that traffic marked with a particular type-of-service bit will not be dropped so long as it remains in profile, and simpler prioritisation-based services.

Towards the right-hand side of the spectrum, flows are typically identifiable in the Internet by the tuple: {source machine, destination machine, source port, destination port, protocol} any of which could be “ANY” (wildcarded).

In the multicast case, the destination is the group, and can be used to provide efficient aggregation.

Flows can be grouped in classes, where each class has an associated service model applied. This can default to best effort.

Through network management, we can imagine establishing classes of long lived flows—enterprise networks (“Intranets”) often enforce traffic policies that distinguish priorities which can be used to discriminate in favor of more important traffic in the event of overload (though in an underloaded network, the effect of such policies will be invisible, and may incur no load/work in routers).

The router service model to provide such classes with different treatment can be as simple as a priority queuing system, or it can be more elaborate.

Although best-effort services can support real-time traffic, classifying real-time traffic separately from non-real-time traffic and giving real-time traffic priority treatment ensures that this traffic sees minimum delays. Non-real-time TCP traffic tends to be elastic in its bandwidth requirements, and will then tend to fill any remaining bandwidth.

We could imagine a future Internet with sufficient capacity to carry all of the world’s telephony traffic. Since this is a relatively modest capacity requirement, it might be simpler to establish “POTS” as a static class which is given some fraction of the capacity overall, and then within the backbone of the network no individual call need be given an allocation (i.e. we would no longer need the call setup/tear down that was needed in the legacy POTS which was only present due to under-provisioning of trunks, and to allow the trunk exchanges the option of call blocking). The vision is of a network that is engineered with capacity for all of the non-best-effort average load sources to send without needing individual reservations.

4.2. Reservations

For flows that may take a significant fraction of the network (i.e. are “special” and can’t just be lumped under a static class), we need a more dynamic way of establishing these classifications. In the short term, this applies to many multimedia calls since the Internet is largely under-provisioned at the time of writing.

RSVP has been standardised for just this purpose. It provides flow identification and classification. Hosts and applications are modified to speak RSVP client language, and routers speak RSVP.

Since most traffic requiring reservations is delivered to groups (e.g. TV), it is natural for the receiver to make the request for a reservation for a flow. This has the added advantage that different receivers can make heterogeneous requests for capacity from the

best effort unsignalled	prioritised by type of service	assured by type of service	assured by aggregate reservation	guaranteed by per-flow reservation
----------------------------	--------------------------------------	----------------------------------	--	--

Fig. 2. Spectrum of internet service types.

same source. Thus with appropriate codecs, RSVP can accommodate monochrome, color and HDTV receivers from a single source (also see Section 5).

Again the routers conspire to deliver the right flows to the right locations.

To support the wildcarding described earlier, RSVP can specify several different filter types. At one extreme, a filter may specify a flow of packets of a specified protocol from a specified port on a sending host to a specified port on the receiving host. Alternatively, shared reservations can be set up, where several specified senders use a single reservation, and at the other extreme, even an unspecified set of senders can share a reservation.

If a network is provisioned such that it has excess capacity for all the real-time flows using it, a simple priority classification ensures that real-time traffic is minimally delayed. However, if a network is insufficiently provisioned for the traffic in a real-time traffic class, then real-time traffic will be queued, and delays and packet loss will result. Thus in an under-provisioned network, either all real-time flows will suffer, or some of them must be given priority.

RSVP provides a mechanism by which an admission control request can be made, and if sufficient capacity remains in the requested traffic class, then a reservation for that capacity can be put in place.

If insufficient capacity remains, the admission request will be refused, but the traffic will still be forwarded with the default service for that traffic's traffic class. In many cases even an admission request that failed at one or more routers can still supply acceptable quality as it may have succeeded in installing a reservation in all the routers that were suffering congestion. This is because other reservations may not be fully utilising their reserved capacity in those routers where the reservation failed.

4.2.1. Billing

If a reservation involves setting aside resources for a flow, this will tie up resources so that other reservations may not succeed, and depending on whether the flow fills the reservation, other traffic is prevented from using the network. Clearly some negative feedback is required in order to prevent pointless reservations from denying service to other users. This feedback is typically in the form of billing.

Billing requires that the user making the reservation is properly authenticated so that the correct user can be charged. Billing for reservations introduces a level of complexity to the Internet that has not typically been experienced with non-reserved traffic, and requires network providers to have reciprocal usage-based billing arrangements for traffic carried between them. It also suggests the use of mechanisms whereby some fraction of the bill for a link reservation can be charged to each of the downstream multicast receivers.

4.3. Differentiated services

Whereas RSVP asks routers to classify packets into classes to achieve a requested quality of services, it is also possible to explicitly mark packets to indicate the type of service required. Of course, there has to be an incentive and mechanisms to ensure that "high-priority" is not set by everyone in all packets, and this incentive is provided by edge-based policing and by buying profiles of higher priority service. In this context, a profile could have many forms, but a typical profile might be a token-bucket filter specifying a mean rate and a bucket size with certain time-of-day restrictions.

This is still an active research area, but the general idea is for a customer to buy from their provider a profile for higher quality service, and the provider polices marked traffic from the site to ensure that the profile is not exceeded. Within a provider's network, routers give preferential services to packets marked with the relevant type-of-service bit. Where providers peer, they arrange for an aggregate higher-quality profile to be provided, and police each other's aggregate if it exceeds the profile. In this way, policing only needs to be performed at the edges to a provider's network on the assumption that within the network there is sufficient capacity to cope with the amount of higher-quality traffic that has been sold. The remainder of the capacity can be filled with regular best-effort traffic.

One big advantage of differentiated services over reservations is that routers do not need to keep per-flow state, or look at source and destination addresses to classify the traffic, and this means that routers can be considerably simpler. Another big advantage is that the billing arrangements for differ-




Fig. 3. Network jitter and packet audio.

entiated services are pairwise between providers at boundaries—at no time does a customer need to negotiate a billing arrangement with each provider in the path³.

5. Transport protocols

So-called real-time delivery of traffic requires little in the way of transport protocol. In particular, real-time traffic that is sent over more than trivial distances is not retransmittable.

With packet multimedia data there is no need for the different media comprising a conference to be carried in the same packets. In fact it simplifies receivers if different media streams are carried in separate flows (i.e., separate transport ports and/or separate multicast groups). This also allows the different media to be given different quality of service. For example, under congestion, a router might preferentially drop video packets over audio packets. In

addition, some sites may not wish to receive all the media flows. For example, a site with a slow access link may be able to participate in a conference using only audio and a whiteboard whereas other sites in the same conference with more capacity may also send and receive video. This can be done because the video can be sent to a different multicast group than the audio and whiteboard. This is first step towards coping with heterogeneity by allowing the receivers to decide how much traffic to receive, and hence allowing a conference to scale more gracefully.

5.1. Receiver adaptation and synchronisation

Best-effort traffic is delayed by queues in routers between the sender and the receivers. Even reserved priority traffic may see small transient queues in routers, and so packets comprising a flow will be delayed for different times. Such delay variance is known as jitter, and is illustrated in Fig. 3.

Real-time applications such as audio and video need to be able to buffer real-time data at the receiver for sufficient time to remove the jitter added by the network and recover the original timing relationships between the media data. In order to know

³ With reservations there may be ways to avoid this too, but they're somewhat more difficult given the more specific nature of a reservation.

how long to buffer for, each packet must carry a timestamp which gives the time at the sender when the data was captured. Note that for audio and video data timing recovery, it is not necessary to know the absolute time that the data was captured at the sender, only the time relative to the other data packets.

As audio and video flows will receive differing jitter and possibly differing quality of service, audio and video that were grabbed at the same time at the sender may not arrive at the receiver at the same time. At the receiver, each flow will need a playout buffer to remove network jitter. Inter-flow synchronisation can be performed by adapting these playout buffers so that samples/frames that originated at the same time are played out at the same time (Fig. 4). This requires that the time base of different flows from the same sender can be related at the receivers, e.g. by making available the times at which each of them was captured relative to a common clock.

5.2. RTP

The transport protocol for real-time flows is RTP [27]. This provides a standard format packet header which gives media specific timestamp data, as well as payload format information and sequence number-

ing amongst other things. RTP is normally carried using UDP. It does not provide or require any connection setup, nor does it provide any enhanced reliability over UDP. For RTP to provide a useful media flow, there must be sufficient capacity in the relevant traffic class to accommodate the traffic. How this capacity is ensured is independent of RTP.

Every original RTP source is identified by a source identifier, and this source id is carried in every packet. RTP allows flows from several sources to be mixed in gateways to provide a single resulting flow. When this happens, each mixed packet contains the source ids of all the contributing sources.

RTP media timestamp units are flow specific—they are in units that are appropriate to the media flow. For example, 8 kHz sampled PCM encoded audio has a timestamp clock rate of 8 kHz. This means that inter-flow synchronisation is not possible from the RTP timestamps alone.

Each RTP flow is supplemented by Real-Time Control Protocol (RTCP) packets. There are a number of different RTCP packet types. RTCP packets provide the relationship between the realtime clock at a sender and the RTP media timestamps so that inter-flow synchronisation can be performed, and they provide textual information to identify a sender in a conference from the source id.




Fig. 4. Inter-media synchronisation.

5.3. Conference membership and reception feedback

IP multicast allows sources to send to a multicast group without being a receiver of that group. However, for many conferencing purposes it is useful to know who is listening to the conference, and whether the media flows are reaching receivers properly. Accurately performing both these tasks restricts the scaling of the conference. IP multicast means that no-one knows the precise membership of a multicast group at a specific time, and this information cannot be discovered, as to try to do so would cause an implosion of messages, many of which would be lost⁴. Instead, RTCP provides approximate membership information through periodic multicast of session messages which, in addition to information about the recipient, also give information about the reception quality at that receiver. RTCP session messages are restricted in rate, so that as a conference grows, the rate of session messages remains constant, and each receiver reports less often. A member of the conference can never know exactly who is present at a particular time from RTCP reports, but does have a good approximation to the conference membership. This is analogous to what happens in a real-world meeting hall; the meeting organisers may have an attendance list, but if people are coming and going all the time, they probably do not know exactly who is in the room at any one moment.

Reception quality information is primarily intended for debugging purposes, as debugging of IP multicast problems is a difficult task. However, it is possible to use reception quality information for rate adaptive senders, although it is not clear whether this information is sufficiently timely to be able to adapt fast enough to transient congestion.

5.4. Scaling issues and heterogeneity

The Internet is very heterogeneous, with link speeds ranging from 14.4 kbit/s up to 1.2 Gbit/s, and very varied levels of congestion. How then can a




Fig. 5. Receiver adaptation using multiple layers and multiple multicast groups.

single multicast source satisfy a large and heterogeneous set of receivers?

In addition to each receiver performing its own adaptation to jitter, if the sender layers [21] its video (or audio) stream, different receivers can choose to receive different amounts of traffic and hence different qualities. To do this the sender must code their video as a base layer (the lowest quality that might be acceptable) and a number of enhancement layers, each of which adds more quality at the expense of more bandwidth. With video, these additional layers might increase the framerate or increase the spatial resolution of the images or both. Each layer is sent to a different multicast group, and receivers can decide individually how many layers to subscribe to. This is illustrated in Fig. 5. Of course, if receivers in a conference are going to respond to congestion in this way, then we also need to arrange that the receivers behind a common bottleneck tend to respond in a synchronised manner. If this is not done, de-synchronised experiments by different receivers will have the net effect that too many layers are always being drawn through a common bottleneck. RLM [20] is one way that this might be achieved, although there is continuing research in this area.

6. Conference control

Conferences come in many shapes and sizes, but there are only really two models for conference control: light-weight sessions and tightly coupled conferencing. For both models, rendezvous mechanisms are needed. Note that the conference control

⁴ Note that a conference policy that restricts conference membership can be implemented using encryption and restricted distribution of encryption keys, of which more later.

model is orthogonal to issues of quality of service and network resource reservation, and it is also orthogonal to the mechanism for discovering the conference [25].

Light-weight sessions are multicast based multimedia conferences that lack explicit conference membership control and explicit conference control mechanisms. Typically a lightweight session consists of a number of many-to-many media streams supported using RTP and RTCP using IP multicast.⁵ Typically, the only conference control information needed during the course of a light-weight session is that distributed in the RTCP session information, i.e. an approximate membership list with some attributes per member.

Tightly coupled conferences may also be multicast based and use RTP and RTCP, but in addition they have an explicit conference membership mechanism and may have an explicit conference control mechanism that provides facilities such as floor control.

The most widely used tightly coupled conference control protocols suitable for Internet use are those belonging to the ITU's H.323 family [17]. However it should be noted that this is inappropriate for large sessions where scaling problems will be introduced by the conference control mechanisms.

In order to try and address this, the ITU is currently standardising H.332, which is essentially a small tightly coupled H.323 conference with a larger lightweight-sessions-style conference listening in as passive participants. It is not yet clear whether H.332 will see large scale acceptance, as its benefits over a simple lightweight session are not terribly obvious. It seems likely that lightweight sessions combined with stream authentication (Section 8.3) might be a more appropriate solution for many potential users.

6.1. Controlling multimedia servers

The Real-Time Stream-control Protocol (RTSP) [28] provides a standard way to remote control a

multimedia server. While primarily aimed at web-based media-on-demand services, RTSP is also well suited to provide VCR-like controls for audio and video streams, and to provide playback and record functionality of RTP data streams. A client can specify that an RTSP server plays a recorded multimedia session into an existing multicast-based conference, or can specify that the server should join the conference and record it.

6.2. Protocols for non-A/V applications

Applications other than audio and video have evolved in Internet conferencing, e.g. Imm [4], Wh [9], NTE [10]. Such applications can be used to substitute for meeting aids in physical conferences (whiteboards, projectors) or replace visual and auditory cues that are lost in teleconferences (e.g., a speaker list application); they also can enable new styles of joint work.

Most non-A/V applications have in common that the application protocol is about establishing and updating a shared state. Loss of information is often not acceptable, so some form of multicast reliability is required. The applications' requirements differ: Some applications make per-participant additions to the shared state that are orthogonal to each other (e.g., whiteboards), some evolve a more closely interrelated common state (e.g., additions to a speaker list must be properly sequenced). Some applications can make use of added bandwidth/react to congestion in an elastic way, others transport data that, although not strictly real-time, is time-critical.

In the IRTF research group on Reliable Multicast [15], work is in progress on common protocol elements that can be used in such applications. At the time of writing, some aspects of reliable multicast are not well-understood, such as the proper way to provide congestion control in a multicast environment. As congestion control is considered an essential element, standards track protocols are not expected before this can be solved.

7. Conference setup

There are two basic forms of conference discovery mechanism. These are session advertisement and

⁵ There is some confusion on the term session, which is sometimes used for a conference and sometimes for a single media stream transported by RTP. In this paper, we prefer to use the less ambiguous term conference except where existing protocols use the term session.

session invitation. Session advertisements are provided using a session directory, and inviting a user to join a session is provided using a session invitation protocol such as SIP [12,26] or H.323 [17].

7.1. Session directories

The rendezvous mechanism for many light-weight sessions is a multicast based session directory. This “broadcasts” session descriptions [11] to all the potential session participants. These session descriptions provide an advertisement that the session will exist, and also provide sufficient information including multicast addresses, ports, media formats and session times so that a receiver of the session description can join the session. The session description protocol (SDP) describes the content and format of a multimedia session, and the session announcement protocol (SAP) is used to distribute it to all potential session recipients.

This mechanism can also be applied to advertised tightly coupled sessions, and only requires that additional information about the mechanism to use to join the session is given. However, as the number of sessions in the session directory grows, we expect that only larger-scale public sessions will be announced in this manner, and smaller, more private sessions will tend to use direct invitation rather than advertisement.

7.2. Session invitation

Not all sessions are advertised, and even those that are advertised may require a mechanism to explicitly invite a user to join a session. Such a mechanism is required regardless of whether the session is a lightweight session or a more tightly coupled session, although the invitation system must specify the mechanism to be used to join the session.

The Session Initiation Protocol (SIP) [12,26] provides a mechanism whereby a user can be invited to participate in a conference. SIP does not care whether the session is already ongoing, or is just being created, and it doesn’t care whether the conference is a small tightly coupled session or a huge broadcast—it merely conveys an invitation to a user in a timely manner, inviting them to participate, and provides enough information for them to be able to

know what sort of session to expect. Thus although SIP can be used to make telephone-style calls, it is by no means restricted to that style of conference.

As users are mobile, it is important that such an invitation mechanism is capable of locating and inviting a user in a location independent manner. Thus user addresses need to be used as a level of indirection rather than routing a call to a specific terminal. The invitation mechanism should also provide for alternative responses, such as leaving a message or being referred to another user, should the invited user be unavailable.

Although SIP is also somewhat lighter weight than H.323 call signalling, the key feature of SIP that distinguishes it from other invitation mechanisms is this support for user mobility. It is explicitly designed to be relayed via proxies, which are typically closer to the caller and so have more up-to-date information than the caller. Several proxies may sometimes relay the call before it reaches the caller’s current location. Without careful design, this could potentially be a security problem, allowing the tracking of a user’s location. SIP is designed from the outset to allow the careful restriction of how location information is released, whilst still allowing the user to be callable.

8. Security

There is a temptation to believe that multicast is inherently less private than unicast communication since the traffic visits so many more places in the network. In fact, this is not the case except with broadcast and prune type multicast routing protocols [5]. However, IP multicast does make it simple for a host to anonymously join a multicast group and receive traffic destined to that group without the other senders’ and receivers’ knowledge. If the application requirement (conference policy) is to communicate between some defined set of users, then strict privacy can only be enforced in any case, through adequate end-to-end encryption.

RTP specifies a standard way to encrypt RTP and RTCP packets using private key encryption schemes such as DES [23]. It also specifies a standard mechanism to manipulate plain text keys using MD5 [24] so that the resulting bit string can be used as a DES

key. This allows simple out-of-band mechanisms such as privacy-enhanced mail to be used for encryption key exchange [19].

8.1. Authentication and key distribution

Key distribution is closely tied to authentication. Conference or session directory keys can be securely distributed using public-key cryptography on a one-to-one basis (by e-mail, a directory service, or by an explicit conference setup mechanism), but this is only as good as the certification mechanism used to certify that a key given by a user is the correct public key for that user. Such certification mechanisms [3] are however not specific to conferencing, and it looks likely that certificates such as those provided by PGP will be most widely used in the near term.

Session keys can be distributed using encrypted Session Descriptions carried in SIP session invitations, or in encrypted session announcements as described below. Neither of these mechanisms provide for changing keys during a session as might be required in some tightly coupled sessions, but they are probably sufficient for many used in the context of lightweight sessions.

Even without privacy requirements in the conference policy, strong authentication of a user is required if making a network reservation results in usage based billing.

8.2. Encrypted session announcements

Session Directories can make encrypted session announcements using private key encryption, and carry the encryption keys to be used for each of the conference media streams in the session. Whilst this does not solve the key distribution problem, it does allow a single conference to be announced more than once to more than one key-group, where each group holds a different session directory key, so that the two groups can be brought together into a single conference without having to know each other's keys.

8.3. Secured "broadcasts"

While private-key encryption is sufficient to exclude non-members from sending or receiving multicast conference traffic, it does mean that all members

of a session are equal. This is normally acceptable for multi-way conferences, but will not be acceptable for many broadcasters who require the ability to ensure that only they can send, perhaps in addition to ensuring that only their paid customers can receive. This is nicely illustrated by the multicast of the Rolling Stones concert in 1994 which was billed as being the first live concert on the Mbone. In fact, this honour goes to a little known band called Severe Tire Damage who had multicast an impromptu concert a year previously. To make their point, just before the Stones were due to go on stage, Severe Tire Damage suddenly started broadcasting one of their songs live to the same multicast group. Clearly commercial broadcasters want to avoid occurrences like this one.

Such secured broadcasts can be performed by encrypting a hash (digitally signing) of each packet with the senders private key of a public-private key pair. The public key is then given to the receivers, and they discard (and prune if possible) any packets that are unsigned. The problem with this is that even encrypting a 128 bit hash with a public key algorithm can be relatively expensive to perform at high packet rates sometimes seen with video. The use of public-key cryptography for this purpose has not yet been standardised, but some such mechanism will clearly be needed before the Mbone becomes an acceptable environment for commercial broadcasters.

9. Summary

This document is an attempt to gather together in one place the set of assumptions behind the design of the Internet Multimedia Conferencing architecture, and the services that are provided to support it.

Fig. 6 shows an example time sequence involved in setting up a light-weight session between two sites. In this case, site A creates a session advertisement, and some time later starts sending a media stream even though there may be no receiver at that time. Some time later, site B joins the session (the multicast routing protocol here is PIM), and starts to receive the traffic. At the earliest opportunity site B also makes an RSVP reservation to ensure the flow quality is satisfactory. This example should be taken as illustrative only—there are different ways to join




Fig. 6. Joining a light-weight multimedia session.

sessions, and different ways to get improved quality of service.

The lightweight sessions model for Internet multimedia conferencing may not be appropriate for all conferences, but for those sessions that do not *require* tightly-coupled conference control, it provides an elegant style of conferencing that scales from two participants to millions of participants. It achieves this scaling by virtue of the way that multicast routing is receiver driven, keeping essential information about receivers local to those receivers. Each new participant only adds state close to them in the network. It also scales by not requiring explicit conference join mechanisms; if everyone were to need to know exactly who is in the session at any time, the scaling would be severely adversely affected. RTCP provides membership information that is accurate when the group is small, and increasingly only a statistical representation of the membership as the group grows. Security is handled through the use of encryption rather than through the control of data distribution.

For those that require tightly coupled conferences, solutions such as H.323 are emerging there too.

There are still many parts of this architecture that are incomplete, and are still the subject of active research. In particular, *differentiated services* for better-than-best-effort service show great promise to provide a more scalable alternative to individual reservations. *Multicast routing* scales well to large groups, but scales less well to large numbers of groups; we expect this will become the subject of significant research over the next few years. *Multicast congestion control* mechanisms are still a research topic, although in the last year several schemes have emerged that show promise. *Layered codecs* show great promise to allow conferences to scale in the face of heterogeneity, but the join and leave mechanisms that allow them to perform receiver-based congestion control are still being examined. We have several working examples of *reliable-multicast-based shared applications*; the next few years should see the start of standardisation work in this area as appropriate multicast congestion control mechanisms emerge. Finally a complete *security architecture* for conferencing would be very desirable; currently we have many parts of the solution, but are still waiting for an appropriate key-distribution architecture to emerge from the security research community.

The Internet Multimedia Conferencing architecture and the Mbone have come a long way from their early beginnings on the DARTnet testbed in 1992. The picture is not yet finished, but it has now taken shape sufficiently that we can see the form it will take. Whether or not the Internet does evolve into the single communications network that is used for most telephone, television, and other person-to-person communication, only time will tell. However, we believe that it is becoming clear that if the industry decides that this should be the case, the Internet should be up to the task.

Acknowledgements

Acknowledgments are due to the End-to-End Research Group, the Int-serv, RSVP, MMUSIC and AVT working groups of the IETF, and discussion

with colleagues at UCL. The earliest clear exposition of some of the ideas here was presented at ACM SIGCOMM 1994 in London by Van Jacobson.

References

- [1] A. Ballardie, P. Francis, J. Crowcroft, An architecture for scalable inter-domain multicast Routing, ACM SIGCOMM 1993, pp. 85–95.
- [2] C. Bormann, Providing Integrated Services over Low-bitrate Links, Internet-Draft draftietf-issll-isslow-03.txt, Work in Progress, March 1998.
- [3] CCITT (Consultative Committee on International Telegraphy and Telephony), Recommendation X.509: The Directory – Authentication Framework, 1988.
- [4] W. Dang, Reliable file transfer in the multicast domain, University of Hawaii Technical Report, August 1993.
- [5] S. Deering, C. Partridge, D. Waitzman, Distance Vector Multicast Routing Protocol, IETF RFC 1075, November 1988.
- [6] S. Deering, Multicast routing in inter-networks and extended LANs, ACM SIGCOMM 88, August 1988, pp. 55–64 and Host Extensions for IP Multicasting, IETF RFC 1112.
- [7] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.-G. Liu, L. Wei, An architecture for wide area multicast routing, ACM SIGCOMM 1994, London, October 1994, Computer Communications Review 24 (4) (1994) 126–135.
- [8] Estrin, Farinacci, Helmy, Thaler, Deering, Handley, Jacobson, Liu, Sharma, Wei, Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification, IETF RFC 2117.
- [9] S. Floyd, V. Jacobson, S. McCanne, C.-G. Liu, L. Zhang, A reliable multicast framework for light-weight sessions and application level framing, ACM SIGCOMM 1995, pp. 342–356.
- [10] M. Handley, J. Crowcroft, Network Text Editor (NTE): A scalable shared text editor for the Mbone, Proc. ACM SIGCOMM 97, Cannes, France, 1997.
- [11] M. Handley, V. Jacobson, SDP: Session Description Protocol, IETF RFC 2327, April 1998.
- [12] M. Handley, H. Schulzrinne, E. Schooler, SIP: Session Initiation Protocol, IETF Internet-Draft, Work in Progress.
- [13] M. Handley, D. Thaler, D. Estrin, The Internet Multicast Address Allocation Architecture, IETF Internet-Draft, Work in Progress.
- [14] V. Hardman, A. Sasse, M. Handley, A. Watson, Reliable audio for use over the Internet, Proc INET '95, Hawaii, Internet Society, Reston, VA, 1995.
- [15] IRTF Research Group on Reliable Multicast, <http://www.east.isi.edu/RMRG/>.
- [16] ITU, Recommendation H.320: Narrow-band Visual Telephone Systems and Terminal Equipment, ITU, Geneva, July 1997.
- [17] ITU, Recommendation H.323: Visual Telephone Systems and Equipment for Local Area Networks which Provide a Non-guaranteed Quality of Service, ITU, Geneva, November 1996.
- [18] V. Jacobson, Congestion avoidance and control, ACM SIGCOMM 88, August 1988.
- [19] J. Linn, Privacy Enhancement for Internet Electronic Mail: Part I: Message Encryption and Authentication Procedures, IETF RFC 1421, February 1993.
- [20] S. McCanne, V. Jacobson, M. Vetterli, Receiver-driven layered multicast, ACM SIGCOMM, Stanford, CA, August 1996, pp. 117–130.
- [21] S. McCanne, M. Vetterli, Joint source/channel coding for multicast packet video, Proc. IEEE Int. Conf. on Image Processing, Washington, DC, October 1995.
- [22] J. Moy, Multicast Extensions to OSPF, IETF RFC 1584, March 1994.
- [23] National Institute of Standards and Technology (NIST), FIPS Publication 46-1: Data Encryption Standard, January 22, 1988.
- [24] R. Rivest, The MD5 Message-Digest Algorithm, RFC 1321, MIT Laboratory for Computer Science and RSA Data Security, April 1992.
- [25] E. Schooler, A distributed architecture for multimedia conference control, ISI Research Report ISI/RR-91-289, November 1991. <ftp://ftp.isi.edu/pub/hpcc-papers/mmcc/mmcc.ps>.
- [26] H. Schulzrinne, Personal mobility for multimedia services in the Internet, IMDS'96, March 4–6, 1996.
- [27] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson RTP: A Transport Protocol for Real-Time Applications, IETF RFC 1889.
- [28] H. Schulzrinne, A. Rao, R. Lanphier, Real Time Streaming Protocol (RTSP), IETF RFC 2326, April 1998.
- [29] K. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoglu, D. Estrin, M. Handley, The MASC/BGMP architecture for inter-domain multicast routing, Proc. ACM SIGCOMM 98, Vancouver, Canada, September 1998.